21

Flux Invariants for Shape

Pavel Dimitrov

Department of Computer Science,
McGill University, Montréal
Québec, Canada

July, 2003

A thesis submitted to the Faculty of Graduate Studies and Research
in partial fulfillment of the requirements of the degree of
Master of Science

List of Algorithms

1 Simple thresholding. 50
2 Generic thinning. 54
3 Discrete skeleton: Thick skeleton algorithm. 57
4 Discrete skeleton: Thin skeleton algorithm. 58
5 Shifting $S_{Thin(d)}(X)$ 65
6 Reordering points after shifting $S_{Thin(d)}(X)$ 66
7 The complete procedure. 69

adjust coordinates such that $\sigma = 1$ and write the following function

$$g(\zeta) = \frac{\cos(\zeta)\,(1+d-\cos(\zeta)) - \sin^2(\zeta)}{\sqrt{(1+d-\cos(\zeta))^2 + \sin^2(\zeta)}}$$

It represents the inner product of normals to the circle and gradient vectors converging to (or emanating from) a single point; see Figure B.4. Thus,

$$f'(s_0) = g'(\zeta_0)$$

as in the figure. However, this holds for a fixed δ. In order to obtain $f''(s)$, it is necessary to let δ change as a function of the boundary parameterization t and, hence, $\delta(s) = \delta(t(s))$. So, assuming $\frac{\mathrm{d}\zeta}{\mathrm{d}s} = 1$,

$$\begin{aligned} f''(s) &= \frac{\mathrm{d}}{\mathrm{d}s} g'(\delta(s), s) \\ &= \frac{\mathrm{d}^2 g}{\mathrm{d}\zeta \mathrm{d}\delta}(\delta, \zeta)\frac{\mathrm{d}\delta}{\mathrm{d}\zeta} + \frac{\mathrm{d}^2 g}{\mathrm{d}\zeta^2}(\zeta) \end{aligned}$$

where

$$\frac{\mathrm{d}\delta}{\mathrm{d}\zeta} = \frac{\mathrm{d}\delta}{\mathrm{d}s} = \frac{\mathrm{d}\delta}{\mathrm{d}t}\frac{\mathrm{d}t}{\mathrm{d}s} = \kappa'(t)\frac{\mathrm{d}t}{\mathrm{d}s}.$$

Thus, $\frac{\mathrm{d}\delta}{\mathrm{d}s} \leq K_2 \frac{R}{d}$. Carrying out the operations we get the following bounds:

$$\begin{aligned} \left|\frac{\mathrm{d}^2}{\mathrm{d}\zeta^2}(\zeta)\right| &\leq \frac{8 + 38d^2 + 25d^3 + 8d^4 + 28d + d^5}{d^5} \\ \frac{\mathrm{d}^2 g}{\mathrm{d}\zeta \mathrm{d}\delta}(\delta, \zeta) &\leq \frac{8 + 16d + 2d^3 + 10d^2}{d^5} \end{aligned}$$

This finishes the proof. □

Bibliography

[1] M. B. Allen and E. L. Isaacson. *Numerical Analysis for Applied Science*. Pure and Applied Mathematics. Wiley-Interscience, 1997.

[2] C. Arcelli and G. S. di Baja. Finding local maxima in a pseudo-euclidean distance transform. *Computer Vision, Graphics, and Image Processing*, 43:361–367, 1988.

[3] C. Arcelli and G. S. di Baja. A one-pass two-operations process to detect the skeletal pixels on the 4-distance transform. *IEEE Transactions on Pattern Recognition and Machine Intelligence*, 11:411–414, April 1989.

[4] J. August, K. Siddiqi, and S. W. Zucker. Contour fragment grouping and shared, simple occluders. *Computer Vision and Image Understanding*, 76(2):146–162, November 1999.

[5] F. Aurenhammer. Voronoi diagram—survey of a fundamental geometric data structure. *ACM Computing Surveys (CSUR)*, 23(3):345–405, 1991.

[6] H. Blum. Biological shape and visual science (part 1). *Journal of Theoretical Biology*, 38:205–287, 1973.

[7] H. Blum and R. N. Nagel. Shape description using weighted symmetric axis features. *Pattern Recognition*, 10(3):167–180, 1978.

[8] S. Bouix. *Medial Surfaces*. PhD thesis, School of Computer Science McGill University, 2003.

[9] J. W. Brandt and V. R. Algazi. Continuous skeleton computation by voronoi diagram. *CVGIP: Image Understanding*, 55(3):329–338, May 1992.

[10] L. Calabi. A study of the skeleton of plane figures. Technical Report 60429, SR-2, Parke Mathematical Laboratories, June 1965.

[11] L. Calabi and W. E. Hartnett. Shape recognition, prairie fires, convex deficiencies and skeletons. *The American Mathematical Monthly*, 75(4):335–342, April 1968.

[12] M. P. D. Carmo. *Differential Geometry of Curves and Surfaces.* Prentice Hall, 1976.

[13] H. I. Choi, S. W. Choi, and H. P. Moon. Mathematical theory of medial axis transform. *Pacific Journal of Mathematics*, 181(1):57–88, November 1997.

[14] J. N. Damon. Global geometry of regions and boundaries via skeletal and medial integrals. *manuscript in preparation*, 2003.

[15] J. N. Damon. Smoothness and geometry of boundaries associated to skeletal structures i. *manuscript in preparation*, 2003.

[16] P.-E. Danielsson. Euclidean distance mapping. *Computer Graphics and Image Processing*, 14(3):227–248, November 1980.

[17] P. Dimitrov, J. Damon, and K. Siddiqi. Flux invariants for shape. In *Computer Vision and Pattern Recognition, 2003*, volume 1, pages 835–841. IEEE Computer Society, June 2003.

[18] P. Dimitrov, C. Phillips, and K. Siddiqi. Robust and efficient skeletal graphs. In *CVPR'00*, volume 1, pages 417–423, 2000.

[19] H. Federer. Curvature measures. *Trans. Amer. Math. Soc.*, 93:418–491, 1959.

[20] S. Fernández-Vidal and G. Malandain. Squelettes euclidiens d'objets discrets n-dimensionnels. Technical Report 2771, INRIA, 2004 route des Lucioles BP 93, 06902 Sophia Antipolis Cedex, France, Janvier 1996.

[21] H. W. Guggenheimer. *Differential Geometry.* Dover, 1977.

[22] M. K. Hu. Visual pattern recognition by moment invariants. *IRE Transaction on Information Theory*, IT-8:179–187, 1962.

[23] J. J. Koenderink. *Solid Shape.* Artificial Intelligence. MIT Press, 1990.

[24] S. G. Krantz and H. R. Parks. Distance to C^k hypersurfaces. *J. Differential Equations*, 40(1):116–120, 1981.

[25] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright. Convergence properties of the nelder-mead simplex algorithm in low dimensions. *SIAM Journal on Optimization*, 9:112–147, 1998.

[26] L. Lam, S.-W. Lee, and C. Y. Suen. Thinning methodologies—a comprehensive survey. *IEEE Transactions on Pattern Recognition and Machine Intelligence*, 14(9):869–885, September 1992.

[27] F. Leymarie and M. Levine. Simulating the grassfire transform using an active contour model. *IEEE Transactions on Pattern Recognition and Machine Intelligence*, 14(1):56–75, January 1992.

[28] M. Leyton. Symmetry-curvature duality. *Computer Vision, Graphics, and Image Processing*, 38:327–341, 1987.

[29] A. Lieutier. Medial axis homotopy. Technical report, LMC/IMAG(Grenoble) and Dassault Systèmes (Aix-en-provence, France, June 2002.

[30] G. Malandain and S. Fernández-Vidal. Euclidean skeletons. *Image and Vision Computing*, 16(5):317–327, April 1998.

[31] J. N. Mather. Distance from a submanifold in euclidean space. *Proceedings of Symposia in Pure Mathematics*, 40(2), 1983.

[32] G. Matheron. *Examples of Topological Properties of Skeletons*, volume 2 of *Image Analysis and Mathematical Morphology*, chapter 11, pages 217–238. Academic Press, 1988.

[33] G. Matheron. *On the Negligibility of the Skeleton and the Absolute Continuity of Erosions*, volume 2 of *Image Analysis and Mathematical Morphology*, chapter 12, pages 239–256. Academic Press, 1988.

[34] F. Meyer. *Skeletons in Digital Spaces*, volume 2 of *Image Analysis and Mathematical Morphology*, chapter 13, pages 257–296. Academic Press, 1988.

[35] F. Mokhtarian. Silhouette-based isolated object recognition through curvature scale space. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(5):539–544, May 1995.

[36] F. Mokhtarian and A. K. Mackworth. A theory of multiscale, curvature-based shape representation for planar curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(8):789–805, August 1992.

[37] F. Mokhtarian and R. Suomela. Robust image corner detection through curvature scale space. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(12):1376–1381, December 1998.

[38] U. Montanari. A method for obtaining skeletons using a quasi-euclidean distance. *Journal of the Association for Computing Machinery*, 15(4), October 1968.

[39] U. Montanari. Continuous skeletons from digitized images. *Journal of the Association for Computing Machinery*, 16(4):534–549, October 1969.

[40] J. R. Munkres. *Topology*. Prentice Hall, 2 edition, December 28 1999.

[41] R. Ogniewicz and M. Ilg. Voronoi skeletons: Theory and applications. In *Proc. of CVPR'92*, pages 63–69, Champaign, IL, June 1992.

[42] R. L. Ogniewicz. *Discrete Voronoi Skeletons*. Hartung-Gorre Verlag, 1993.

[43] R. L. Ogniewicz and O. Kübler. Hierarchic voronoi skeletons. *Pattern Recognition*, 28:343–359, 1995.

[44] J. O'Rourke. *Computational Geometry in C*. Cambridge University Press, second edition, September 1998.

[45] T. Pavlidis. *Algorithms for Graphics and Image Processing*. Computer Science Press, 1982.

[46] M. Pelillo, K. Siddiqi, and S. W. Zucker. Matching hierarchical structures using association graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(11):1105–1120, 1999.

[47] A. Rosenfeld and J. L. Pfaltz. Sequential operations in digital picture processing. *Journal of the ACM*, 13:471–494, October 1966.

[48] K. Siddiqi, S. Bouix, A. Tannenbaum, and S. W. Zucker. Hamilton-jacobi skeletons. *International Journal of Computer Vision*, 48(3):215–231, 2002.

[49] K. Siddiqi, A. Shokoufandeh, S. J. Dickinson, and S. W. Zucker. Shock graphs and shape matching. *International Journal of Computer Vision*, 35(1):13–32, 1999.

[50] M. Singh, G. Seyranian, and D. Hoffman. Parsing silhouettes: The short-cut rule. *Perception and Psychophysics*, 61:636–660, 1999.

[51] N. J. C. Starchan, P. Nesvadba, and A. R. Allen. Fish species recognition by shape analysis of images. *Pattern Recognition*, 23(5):539–544, 1990.

[52] H. Talbot and L. Vincent. Euclidean skeletons and conditional bisectors. In *Visual Communications and Image Processing*, volume 1818 of *Proc. SPIE*, pages 862–876, Boston MA, November 1992. SPIE.

[53] H. Tek and B. B. Kimia. Symmetry maps of free-form curve segments via wave propagation. *IJCV*, to appear 2003.

[54] L. Vincent. Efficient computation of various types of skeletons. In *Medical Imaging V*, volume 1445 of *Proc. SPIE*, pages 297–311, San Jose CA, February 1991.

[55] F. W. Warner. *Foundations of Differentiable Manifolds and Lie Groups*, volume 94 of *Graduate Texts in Mathematics*. Springer-Verlag New York, Incorporated, first edition, January 1983.

[56] Y. Xia. Skeletonization via the realization of the fire front's transform propagation and extinction in digital binary shapes. *IEEE Transactions on Pattern Recognition and Machine Intelligence*, 11(10):1076–1086, October 1989.

Chapter 1

Introduction

Despite its obvious importance and ubiquitous existence, we have yet to determine what visual shape really is. Informal conversations frequently refer to objects by describing aspects of their visual form and, as such, shape is only attributed subordinate meaning, one that may change as circumstance dictates. Yet, regardless of context, we seem to be able to *remember* shape, as though it is an entity in its own right. Thus, from a philosophical standpoint, it is unsatisfactory to be limited to an indirect understanding of visual form.

The exercise of assigning meaning to visual shape is not purely academic in nature, however. The field of computational vision would indeed reach a new level if a reasonably general solution were found: image segmentation, object identification and recognition are just a few of the major areas such a discovery would have an impact on. The higher-level applications range from helping the visually impaired read, to understanding the behavior of certain illnesses, to exploring dangerous sites with autonomous robots. The seed of practical interest, however, is the problem of finding a good representation for shape.

This thesis focuses on a particular class of representations of visual form, those based on the Blum skeleton [6]. We shall argue that such a representation has a number of desirable properties, which any good shape descriptor should possess. We shall begin this discussion in Section 1.1 but will keep it on a relatively high level leaving the details to Chapter 2. We shall then, in Section 1.2, review some of the previous work on skeletonization and identify the major computational difficulties encountered in the past. Finally, in Section 1.3, we shall briefly discuss how the method presented in this thesis deals with those issues and we shall identify the contributions of this work.

1.1 Shape and Shape Representation

Koenderink (see [23, p. 15]) stipulates that visual shape must be given an operational definition; that is, one that may possibly change from task to task. "Thus," he writes, "things do not 'have a shape' the way Santa Claus has a red suit," which, from a practical standpoint, suggests that the same measurement may *have to* be interpreted differently under different circumstances. This is in accordance with Blum's early remarks (see [6]) on the problems pertaining to biological shape. He identifies two types of problems:

- The first is the need to define a "taxonomy," that is, a hierarchically structured representation of shape where subgroups of shapes arise naturally. In particular, it is desirable that visual form be described by constituent subforms. This is not only intuitively satisfactory—since we seem to deal with information in a highly structured (hierarchical) manner—but computationally as well; for example, the problem of matching shapes is thus broken into two subproblems: matching parts and combining this information to test for group membership. There is also a hint at efficiency since, if the hierarchy is a wide tree, then searching may be as fast as $\log n$ where n is the number of parts.

- The second problem is concerned with how organisms perceive and organize the visual form of other organisms. In particular, it is important to be able to tell different organisms apart as well as recognize similar ones. Thus, looking at a panther, the observer should perceive the same visual shape regardless of his relative position to the cat (translation) or whether his head is tilted to one side or not (rotation); however, he should also be able to realize that the gazelle beside the panther has a different visual shape.

A machine representation of shape should incorporate the above observations, but it should also provide the means for an effective and intuitive communication. Below we list a number of criteria that a good shape representation should satisfy:

1. **Completeness.** It should be able to represent a large class of shapes. The largest class is the bounded sets in $\mathbb{R}^2$.

2. **Hierarchy.** It should allow for an easy extraction of parts, or sub-shapes.

3. **Invariance.** In general, the shape itself should remain the same even if some Euclidean transformation is performed on the input. Hence, rotations, translations and scalings should not have an impact on the percept in many situations, but, under certain circumstances, this need not hold. For example, a horizontal bar (a twig, say) and a vertical one may be important higher level organizational processes. Consequently, the representation should allow for an interpretation invariant under the aforementioned transformations, but it should also keep absolute information.

4. **Noise.** It should be able to deal with noise; in other words, there should be a natural interpretation which emphasizes differences between major features and minor variability.

5. **Metric.** There should be a way to define, at least locally, a pseudo-metric (although a metric might be necessary for some applications) thereby defining clusters of "alike" shapes.

6. **Language.** Ideally, it should provide the means to extend a human language in order to allow for a more efficient communication with machines.

There are at least two classes of direct shape representation in the computer vision literature: boundary based and region based. In the former case, the completeness and invariance criteria are satisfied and there are several approaches that define a shape similarity metric.[1] However, small variations are captured and some extrinsic regularization (E.g. Gaussian smoothing) must be applied to deal with noise. Further, there is no intuitive notion of parts (see August *et al.* [4]), even though several efforts[2] break the contour into segments each of which is meant to describe portions of the shape.

On the other hand, the region based approaches emphasize a more intuitive notion of sub-shapes as well as the more global symmetries. An advantage these approaches have over boundary based ones, is that they encode the topological information of the shape as well as inter-point relations (such as distance, or relative distance, between two points in a shape). However, region based methods may have large memory requirements.

[1]For example, the work by Mokhtarian and collaborators [36, 35, 37] on curvature scale space. Also, see Hu [22] for an idea based on invariant moments.

[2]e.g. Singh *et al.* [50] from a human perception perspective and Starchan *et al.* [51] from a practical standpoint

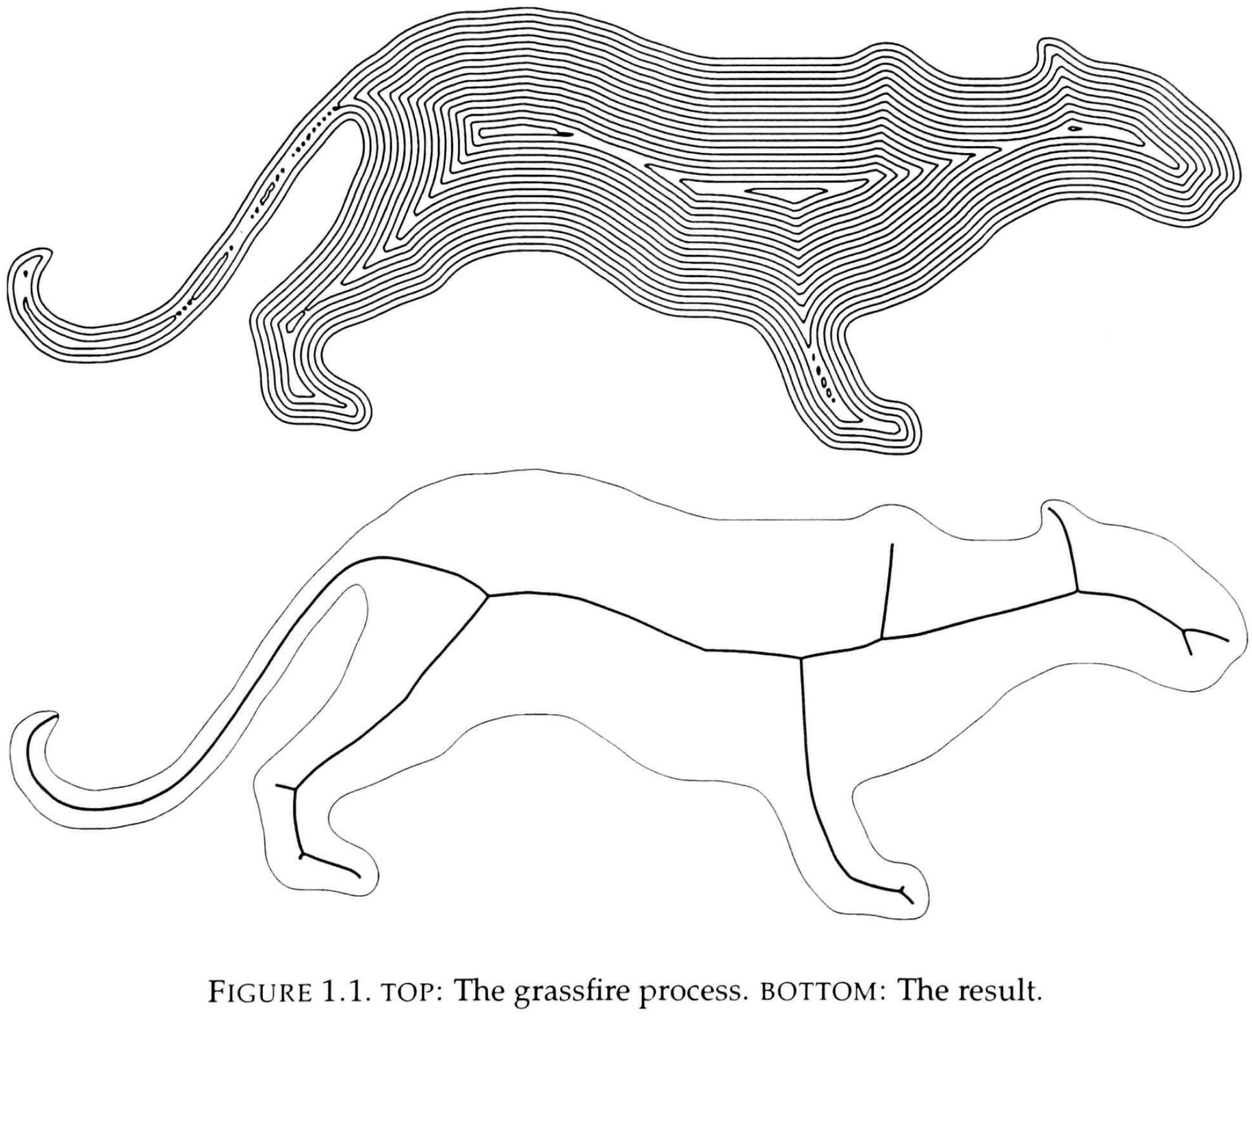

FIGURE 1.1. TOP: The grassfire process. BOTTOM: The result.

A shape representation which combines the advantages of these two general approaches was first introduced by Blum in [6]. He called it the "skeleton" of a shape and it is perhaps best explained through his well known *grassfire* analogy. Imagine that the shape is given as a perfectly dry and flat grass region surrounded by a wet area. The boundary is then set on fire everywhere simultaneously and the front advances inward at a constant speed (see Figure 1.1 top). As fire fronts meet, quench points form and, eventually, the fire is extinguished. The set of quench points sites is a sort of "stick-figure" of the original which is also called the *skeleton* or *medial axis* of the shape. If one keeps the time of formation of the quench points (distance to the boundary), the resulting structure is called the *medial axis transform* (MAT) because complete information about the process is retained; in particular, the original shape may be regenerated by an inverse grassfire.

The MAT is a good representation for visual form: it is easily seen that the grassfire process defines a skeleton for all bounded shapes; in Chapter 2, we shall prove that the medial axis can be realized as a graph which satisfies the invariance criterion above; we shall also see, as a consequence of results from Chapter 4, that there is a natural way to deal with certain types of noise in the input shape; and there is a large selection of shape metrics based on the skeleton (e.g. Pellilo *et al.* [46] and Siddiqi *et al.* [49]). Consequently, significant effort has been made to compute the MAT but, as the next section demonstrates, the problem is far from trivial and progress has been slow.

1.2 Computing Skeletons

The past three decades have produced an enormous number of algorithms that try to compute the skeleton (or a similar representation) of a 2D visual shape. The approaches are so varied that a general classification doing justice to all is beyond the scope of this thesis. Therefore, rather than provide a complete overview, we shall concentrate in this section on the methodologies closely related to the class of algorithms presented in this thesis. Four groups are discussed: topological thinning, distance map methods, grassfire simulation and wave propagation, and Voronoi Diagram (VD) approaches.

1.2.1 Topological Thinning

One of the properties of the skeleton (as shown in Chapter 2) is that it is topologically equivalent to the original shape. In practice, this means that both will have the same number of connected components and holes, and there is a natural and unique way of mapping components (or holes) in the shape to components (or holes) in the skeleton. Topological thinning approaches are characterized by ensuring that the result of skeletonization will have this property. All such algorithms assume that the shape is somehow given on a discrete grid and remove pixels from the boundary according to some criteria. The methods in this category can be further subdivided into two groups:

- *Thinning by contour peeling.* The main idea is to discretely simulate the grass-fire process by successively removing layers of pixels from the shape's boundary. Lam *et al.* present a thorough overview in [26] covering both sequential and parallel implementations. In Chapter 9 of [45], Pavlidis surveys some of the earlier approaches and the issues they present are made clear.

  Although most of these algorithms do get a "skeleton" as a result, the term is a misnomer in the sense of Blum. Blum's definition may be seen as given on $\mathbb{R}^2$, that is, in the continuous case; hence, for unrestricted shapes discrete algorithms may only *approximate* the MAT, not compute it. Therefore, the approximations obtained in this manner do not necessarily have the desired properties (as listed in the previous section). In particular, they are not very consistent under rotations of the shape.

- *Augmented approach.* Vincent [54] is the best representative of this method. The idea is to use knowledge from other approaches to define "anchor points" (i.e. that cannot be removed by the thinning procedure) and then apply a thinning technique which removes boundary points in some order while making sure the topology is left unchanged. The result of this strategy is a topologically equivalent representation of the original shape, but may not be thin if the set of anchor points is not. For example, Talbot and Vincent [52] use local speed of creation of quench points as a determinant for anchor points, but they approximate this value and the thresholding cannot be shown to give a thin anchor set in general. Malandain and Vidal [30, 20] use this criterion as well as the distance to boundary to extract the anchor points but fail to demonstrate how a thin skeleton can be obtained for all input shapes.

1.2.2 Distance Map Methods

An equivalent definition for the skeleton of a shape (see Chapter 2) may be given by using the Euclidean metric and the notion of a maximally inscribed disk. A disk is maximally inscribed in a shape if there is no other disk inside the shape which completely contains it. Thus, the skeleton consists of all points in the shape which are centers of maximally inscribed disks. The role of the metric here is to define what the shape of the disk is, i.e. the Euclidean metric gives a circle, the Manhattan distance gives a square, and so on. Therefore, one can generalize the notion of skeleton by simply leaving the choice of metric open. Computing the skeleton, then, is just a matter of finding the centers of those disks which can be done with the following observation: Centers of maximal disks are local maxima of the distance function.

It turns out that skeletonization algorithms based on non-Euclidean metrics (e.g. [47, 38] and, more recently, [3]) are often much faster than those based on Euclidean metrics. Further, if the shape is given on a grid, the *city block* metric for example, makes the skeleton lie on the same grid. However, a significant drawback of this representation is that, like the result of pure thinning methods, it is sensitive to rotations and is hence unstable. The problem stems from the metric itself, because methods based on quasi-Euclidean and Euclidean distances are better behaved in practice and known to be stable in theory.

Arcelli and Sanniti di Baja [2] and Danielsson [16] are two examples of algorithms based on the local maximum observation for pseudo-Euclidean and Euclidean metrics, respectively. The resulting representations of shape are stable, but may be thick and possibly disconnected.

Finally, Montanari [39] suggests a completely different skeletonization approach using the distance function. Instead of using the discrete grid to approximate the skeleton, Montanari develops an analytic simulation which computes the exact structure (for the Euclidean metric) that Blum describes. The crucial assumption, however, is that the shape can be accurately described by a polygon. Unfortunately, this does not always hold because noise in the input data can drastically change the polygonal representation. Hence, this exact algorithm is of limited scope.

1.2.3 Grassfire Simulation and Wave Propagation

The grassfire analogy used to present the skeleton not only gives a definition, but it also suggests an algorithm for computing the medial axis. Pure thinning on a discrete grid was just a first attempt at simulating the process and was not guaranteed to give the best approximation to the skeleton. The main problem stems from the fact that the simulated fire front does not propagate homogeneously. Xia [56] proposes a modified thinning strategy which addresses this issue but the metrics used are non-Euclidean.

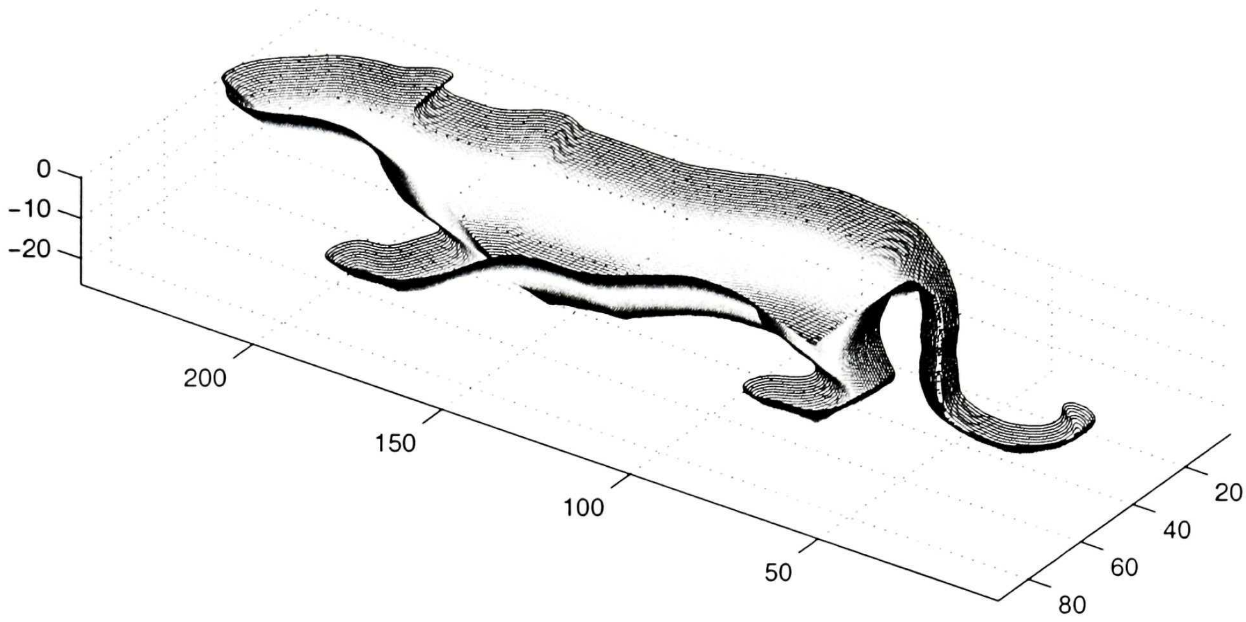

FIGURE 1.2. 3D surface under image plane induced by the distance function.

A more sophisticated approach employing active contours (*snakes*) was developed by Leymarie and Levine in [27]. A snake is defined as a model of a deformable curve composed of abstract elastic materials: strings and rods; thus, it can stretch and bend in a controllable fashion. A grassfire is simulated by letting a snake fall down a surface defined by the Euclidean distance function under the image plane; see Figure 1.2. The idea is indeed interesting and can circumvent many of the shortcomings from previous algorithms, but several sensitive operations must be tuned to produce the desired result; for example, critical points (curvature extrema in particular) of the boundary must be extracted which, on a discrete grid, may be as good as noise if the extraction process does not pick a large enough window.

Another interesting approach for simulating the grassfire is presented in Tek and Kimia [53]. Their method is inspired by numerical algorithms for solving PDE's and the observation that the Eikonal equation can be used to describe isotropic propagation of fire. The shapes are assumed to be given on a discrete grid which they refer to as the fixed image grid. Discrete waves are propagated on it but shocks are detected on a dynamic grid formed by bisectors of the line segments that describe the wave front.

1.2.4 Voronoi Diagram Approaches

It was first observed by Kirkpatrick, over twenty years ago, that the skeleton of a polygonal shape is a subgraph of the Voronoi Diagram—a partitioning of the plane based on proximity to individual line segments (see Aurenhammer [5] for a survey on VD's). This relationship is important because VD's may be computed very efficiently, in $O(n \log n)$ time[3] for a polygon of n vertices, however, as observed previously, polygonal approximation of shapes is quite unstable.

On the other hand, a discrete point sampling of a shape's contour may be a better alternative. Brandt and Algazi [9] show that the skeleton can be well approximated by the Voronoi diagram of such a set. Specifically, they prove that, as the maximum gap between consecutive sample points tends to zero, the Voronoi diagram tends to the skeleton of the shape. However, for any finite sampling, the VD will contain many edges which may not be part of the skeleton. Consequently, any practical algorithm based on their observation must perform some sort of pruning—diagram simplification based on leaf node removal—in order to obtain the approximation of the medial axis.

The solution suggested by Brandt and Algazi [9] is a heuristic approach based on the fitness of regeneration: branches are removed if the regenerated boundary is within some tolerance of the original shape. The combinatorial complexity of the representation after pruning is not easy to determine; in particular, there is no notion of minimal complexity for a given tolerance.

Ogniewicz [42], Ogniewicz and Ilg [41], and Ogniewicz and Kübler [43] provide other heuristics for pruning the Voronoi diagram. Their approaches define the so-called "residual functions" that measure the importance of Voronoi edges. However, as in the previous pruning scheme, there is no real notion of a minimal representational complexity given a threshold.

[3] See O'Rourke [44] for algorithms and analysis.

Abstract

The medial axis of a 2D shape consists of those points in the plane for which the signed distance function induced by the shape is singular. Also known as the skeleton, this structure is particularly well suited to represent visual form: it is thin, homotopic to the original shape, invariant under rigid transformations of the plane, can be used to completely recover the shape, and, most importantly, it has a natural interpretation as a graph. However, despite its virtues, an accurate and robust procedure to compute the medial axis has proved difficult to develop.

In this thesis, we present a solution to this problem in two steps. First, we show that the average outward flux through a region shrinking to a point of the distance function gradient field has a limit everywhere on the plane and is nonzero only at skeletal points. Then, we develop an algorithm based on the discrete interpretation of this behavior and obtain a skeletonization procedure which is essentially parameter free.

Overall, Voronoi diagram approaches compute a connected, i.e. topologically equivalent to the initial shape, approximation of the skeleton which, in theory, may be as close as the sampling permits, but can fail to provide consistent representations. The main problem is that pruning may yield significantly different combinatorial structures, even when the starting point sets are very similar.

1.2.5 Computational Difficulties

The overview in the present section discussed several techniques used to compute the skeleton of a visual shape, each with its own set of advantages and shortcomings. The following is a list, based on this discussion, that collects the computational difficulties commonly encountered in skeletonization algorithms:

- It is a challenge to compute a skeleton approximation that is stable under rotations.
- The output may be thick.
- The topology of the shape may be different from that of the computed skeleton.
- The output may be sensitive to small changes in the input shape.
- Parameters may have to be tuned in an *ad hoc* fashion.

1.3 Contributions

In this thesis, we present a skeletonization algorithm that overcomes the computational difficulties identified in Section 1.2.5. We develop a tool—the average outward flux—which is used to classify points in the plane $\mathbb{R}^2$ as either medial or non-medial. Then, we show how this criterion may be approximated on a discrete lattice and used in a discrete skeletonization algorithm. We prove that the result of this procedure will give a thin structure (only two neighbors in the 3x3 neighborhood of a point) which is no more than a lattice spacing away from the real skeleton. The points in this approximation can then be shifted toward the medial axis until they are no more than a user-specified tolerance from it. Finally, we derive the conditions that guarantee this algorithm to give the desired result which also demonstrates that, if the shapes are originally given on a gird, the procedure only needs the tolerance to be specified. Specifically,

- We discuss the notion of planar shapes and their skeletons and provide an account of the differential and combinatorial properties of these and associated structures (e.g. induced distance function); see Chapter 2.
- We introduce α-skeletons as simplifications of the complete skeleton and show how the two are related (Theorem 2.31).
- We reformulate the Divergence Theorem and use it to establish a criterion that separates skeleton points from non skeleton points in the plane $\mathbb{R}^2$. (Chapter 3) Specifically, the limiting behavior of the *average outward flux* of the gradient of the distance function is described and shown to be strictly negative on skeletal points and only on such points.
- The case of the average outward flux through shrinking circular neighborhoods (Section 3.4) is analyzed in full. We show that, at almost all skeletal points $P \in \mathbb{R}^2$ (all but finitely many), this value is the sine of the object angle; see Table 3.1.
- We translate the continuous behavior into a discrete algorithm (Chapter 4). We analyze the numerical approximations to the average outward flux and estimate, by providing bounds, the error. We thus show how parameters should be chosen in order to obtain a skeleton approximation no further than a specified tolerance from the real object (Theorem 4.8).
- The algorithm we suggest computes arbitrarily close approximations to the α-skeleton of a shape where $\alpha \geq 36°$ (see Theorem 2.31 and Section 4.6).

Organization

The rest of this thesis is organized in three chapters. We shall give shape and its skeleton a formal definition in Chapter 2, and present many of the ensuing properties of these structures. Then, in Chapter 3, we shall extend the Divergence Theorem and derive the criterion for identifying skeletal points. Finally, in Chapter 4, we shall discuss the issues involved in using this criterion on a discrete lattice and develop a skeletonization algorithm based on that analysis.

Chapter 2

Definitions and Properties

Before trying to compute the skeleton of a shape, we must introduce these concepts formally. In this chapter, we shall give precise definitions for both shape and skeleton and, in the process, we shall also introduce the vast majority of concepts needed throughout this thesis.

We'll begin, in Section 2.1 by presenting and briefly discussing the definition of shape. We shall then examine the differential properties of the signed distance function induced by a shape (Section 2.2) and we'll see how skeletons may be defined as singularities of this function in Section 2.3; several of its properties will be derived there as well. Finally, we shall explore the adequacy of the skeleton (transform) as a representation for shape in Section 2.4 and we'll introduce α-skeletons.

2.1 Shapes

We begin, in this section, by quickly revisiting standard notions from point-set topology needed for a formal discussion of shape. Once the necessary notation is established and the required definitions are in place, shape will be given mathematical meaning which shall be the basis for our subsequent analysis.

Our investigation will be restricted to the real plane $\mathbb{R}^2$ and we shall assume the standard Euclidean metric and induced topology. The next four definitions establish the notations and introduce the concept of boundary for subsets of $\mathbb{R}^2$.

Definition 2.1. The *Euclidean norm,* denoted $\|\cdot\| : \mathbb{R}^n \to \mathbb{R}$ is defined as

$$\|\mathbf{v}\| = \sqrt{\sum_{i=1}^{n} v_i^2}$$

where $\mathbf{v} = (v_1, \ldots, v_n)$. The *Euclidean metric* $\mathrm{d}(P,\ Q) : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ is

$$\mathrm{d}(P,\ Q) = \|P - Q\| .$$

We shall also denote by $\mathrm{d}(P,\ S)$ the closest distance from a point P to the set $S \subset \mathbb{R}^2$, i.e.

$$\mathrm{d}(P,\ S) = \inf_{Q \in S} \mathrm{d}(P,\ Q) .$$

Definition 2.2. Let $P \in \mathbb{R}^n$ and $r > 0$. The n-dimensional *open ball* is

$$\mathcal{B}_r(P) = \{Q \in \mathbb{R}^n \ : \ \mathrm{d}(P,\ Q) < |r|\}$$

A set $X \subset \mathbb{R}^2$ will be called *bounded* if and only if there is a $P \in \mathbb{R}^2$ and a large enough $r > 0$ such that $X \subset \mathcal{B}_r(P)$.

Definition 2.3. The *interior* of a set $X \subseteq \mathbb{R}^n$ is

$$\mathrm{int}(X) = \{P \in X \ : \ \exists \varepsilon > 0,\ \mathcal{B}_\varepsilon(P) \subseteq X\} .$$

A *limit point* $Q \in \mathbb{R}^n$ of X if and only if for any $\varepsilon > 0$, $\mathcal{B}_\varepsilon(Q) \cap \mathrm{int}(X) \neq \varnothing$. The collection of all limit points of X is called the *boundary* of X and is denoted ∂X. The *closure* of X is given by

$$\overline{X} = \mathrm{int}(X) \cup \partial X .$$

Definition 2.4. A set $X \subseteq \mathbb{R}^n$ is *open* if $X = \mathrm{int}(X)$. A set X is *closed* in $\mathbb{R}^n$ if and only if $X = \overline{X}$.

In order to talk about the boundary of a shape, we need to establish the notion of real analytic curves. As the following definition illustrates, these are a special kind of differentiable curves.

Definition 2.5. Let $f : \mathbb{R} \to \mathbb{R}$ be a function such that the k-th derivative exists and is continuous, then $f \in C^k$. Any function $f \in C^\infty$ must have continuous derivatives of all orders, but f need not have a converging Taylor series on all points in its domain of definition. A C^∞ function is called *real analytic* and belongs to C^ω if it can be written as a convergent Taylor series for all points in its domain.

Now, if $\gamma : \mathbb{R} \to \mathbb{R}^2$ is a curve defined by $\gamma(t) = (x(t), y(t))$ where both $x(t)$ and $y(t)$ are C^k, then the curve itself will be referred to as being C^k as well. Similarly, the curve is called *real analytic* if and only if both $x(t)$ and $y(t)$ are real analytic functions.

Remark 2.1. If $k < l$ are integers, then if a function is C^l, it must also be C^k. However, the converse need not be true.

Definition 2.6 (Shape). A 2D *shape* X has two defining characteristics:

i. it is the closure of a bounded connected open subset in $\mathbb{R}^2$, i.e. $X = \overline{\mathrm{int}\,(X)}$;

ii. the boundary, ∂X, of X consists of a finite number of mutually disjoint closed curves. Each boundary curve does not self-intersect and is composed of finitely many real analytic curve segments.

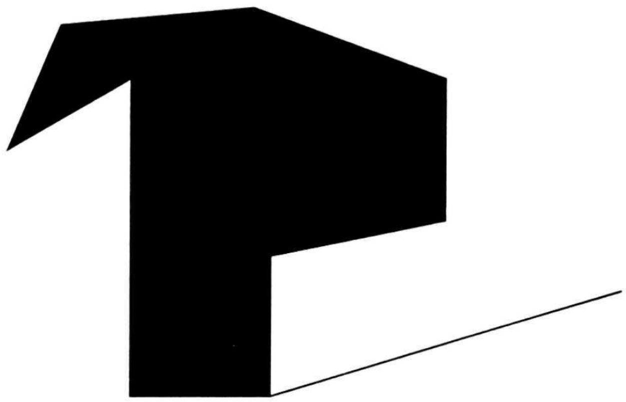

FIGURE 2.1. A disallowed shape. The line connected to the base is a portion of the shape that has no thickness. The interior of this shape only consists of the interior of the "P"; hence, the closure of the interior would be the same shape, with the line removed.

This definition for shape encompasses the largest class of objects which are well behaved mathematically and afford a reasonable representation. Let us examine the first condition. If it is violated, then the object on Figure 2.1 would be valid. There, a piece of the boundary—the line connected to the base of the shape—is part of the figure but has no real thickness. It could be argued that such objects do indeed exist, however our sensing technologies (e.g. digital cameras) do not, at present, allow to detect infinitely thin strings and to correctly identify them as such. Therefore, allowing shapes as that on Figure 2.1 only adds to the representational burden without including any additional observable phenomena.

There are more practical considerations that lead to this definition for shape, but we must postpone this discussion until after the skeleton of a shape is formally introduced in Section 2.3. Thus, the next section presents the necessary preliminaries, namely, the notion of a distance function induced by a shape.

2.2 Euclidean Distance Function

In order to define the skeleton of a shape X, we need to establish several results concerning the induced signed distance function D_X. This section will discuss the differentiability of D_X by analyzing its behavior with respect to the shape.

However, before anything can be said about D_X, we must define it. The signed distance function D_X is Euclidean distance to the boundary of X which is negative for points inside the shape and positive otherwise. Formally,

Definition 2.7. Given a shape X, the signed *distance function* $D_X : \mathbb{R}^2 \to \mathbb{R}$ is defined as

$$D_X(P) = \chi_X(P) \inf_{Q \in \partial X} d(P, Q)$$

where $d(P, Q)$ is the Euclidean distance between $P \in \mathbb{R}^2$ and $Q \in \mathbb{R}^2$ and $\chi_X : \mathbb{R}^2 \to \{-1, 1\}$ is

$$\chi_X(P) = \begin{cases} -1 & \text{if } P \in X \\ 1 & \text{if } P \notin X \end{cases}. \tag{2.1}$$

It turns out that, as a differentiable function, D_X behaves in essentially the same way as the boundary of the shape; but to make this precise, we need to introduce the *boundary support* for a planar point.

Definition 2.8. Let X be a shape and $P \in \mathbb{R}^2$ a point. The *boundary support* of P is the set of closest boundary points to P, i.e.

$$P_C = \{Q \in \partial X \mid \mathrm{d}(P, Q) = \mathrm{d}(P, \partial X)\}.$$

Theorem 2.9. *If $|P_C| = 1$, let $Q \in \partial X$ be the unique point in P_C. If ∂X is piecewise C^k, then*

(a) If $P \notin \partial X$, then

$$\nabla D_X(P) = \chi_X(P) \frac{Q - P}{\|Q - P\|} \tag{2.2}$$

where $Q - P$ denotes the vector from P to $Q \in P_C$ and $\chi_X(P)$ is as in Equation 2.1.

(b) If ∂X is C^k, with $k \geq 1$, near Q, then $D_X(P)$ is C^k near P.

The first part of the theorem is due to Federer [19, 4.8(3)] and the second is a direct consequence of a result by Krantz and Parks [24].[1] Note that, even though part (b) holds for an arbitrary $k \geq 1$, the assumptions on the shapes we consider here imply that $k \geq 1$ actually means $k = \infty$. Similarly, the next Corollary (shown in Appendix A), only needs the boundary to be piecewise C^k.

Corollary 2.10. *Let $P \in \mathbb{R}^2 - \partial X$. Assume $|P_C| = 1$ and let $Q \in \partial X$ be the unique point in P_C. If ∂X is C^0 near Q, then $D_X(P)$ is at least C^1 near P.*

Theorem 2.9 and Corollary 2.10 reveal much of the behavior of the vector field given by the gradient of the distance function. Indeed, they show that ∇D_X has magnitude one and is continuous on all interior points (i.e. $P \in \operatorname{int}(X)$) which have a trivial boundary support. Thus, ∇D_X can be seen as a real-valued function there—only the orientation of the vector is relevant. The significance of this fact will be made explicit in the next section where we shall see that almost all points in the plane have trivial boundary support.

2.3 Skeletons and Skeleton Transform

Now that shape is formally defined and the induced distance function is in place, we can finally turn our attention to the skeleton. We shall begin this section by presenting two definitions of such a structure and will then show that they are equivalent. This will be useful in studying the medial axis as it will allow us to see the structure in two different ways, each contributing to the list of properties presented here. The interplay of these formulations will allow us to show that the skeleton is a thin structure which is topologically equivalent to the original shape. This will enable us to classify skeletal points so as to suggest a natural combinatorial structure—a finite graph—on the medial axis.

We begin, then, by recalling Blum's definition of skeleton based on the grass-fire analogy. Imagine an ideally homogeneous grass field, perfectly flat and where there is no wind. Suppose now that a perfectly dry and ready to burn bounded region is surrounded by wet grass (which cannot catch on fire) and ignite the border of this region, all at once. The fire will then propagate inward at a constant speed,

[1] Krantz and Parks [24] show the claim for an arbitrary (but finite) dimension manifold. Mather [31], on the other hand, provides even more information about the signed distance function by providing normal forms; he has a list for manifolds of up to dimension 7. Yet another source with similar results is Matheron [33] where differential properties of the distance function are shown with the skeleton in mind.

occasionally extinguishing itself at the sites where fronts meet until the whole region is burnt. The quench points are what Blum called the "skeleton" of the shape (the dry region).

Observe that, at all such points, one can draw a disk with radius equal to the closest distance to the boundary and this disk will be maximal; that is, there will be no other disk which strictly contains the first one and is itself completely inside the shape. Formally, this translates into the following definition:

Definition 2.11. A ball $\mathcal{B}_r(P)$ is said to be *inscribed* in a shape X whenever $\mathcal{B}_r(P) \subseteq X$. A ball $\mathcal{B}_r(P)$ is *maximally inscribed* if and only if there is no inscribed ball $\mathcal{B}_{r'}(Q)$ which completely contains $\mathcal{B}_r(P)$ and is strictly larger, i.e. $r' > r$.

Intuitively, it should be clear that the grass fire formulation above is equivalent to defining the medial axis of a shape in the following manner:

Definition 2.12. The *medial axis*, $\mathbf{MA}(X)$, of a shape X is the locus of maximally inscribed disks; that is, a point $P \in X$ is in $\mathbf{MA}(X)$ if and only if there exists a maximally inscribed ball in X centered at P.

We shall not show that equivalence here because Blum's definition is not particularly useful to our efforts in this chapter (or the rest of the thesis), but it should be pointed out that the report [10] by Calabi contains a proof.[2] Definition 2.12, on the other hand, will prove very useful, especially in conjunction with the following one.

Definition 2.13. The *skeleton* of a shape X, denoted $\mathrm{Sk}(X)$, is the set of points in X that have more than one closest point to the boundary of the shape; formally

$$\mathrm{Sk}(X) = \{P \in X \ : \ |P_C| \geq 2\}.$$

The *skeleton transform*, denoted $\mathbf{ST}(X)$, is $\mathbf{ST}(X) = \mathrm{Sk}(X) \times R$ (for $R \subset \mathbb{R}_+ - \{0\}$) where $(P, r) \in \mathbf{ST}(X)$ if and only if $P \in \mathrm{Sk}(X)$ and $r = |D_X(P)|$.

Remark 2.2. Contrast Definition 2.13 with Theorem 2.9 and Corollary 2.10. The skeletal points are exactly the locations in the interior of the shape where D_X is not smooth; in fact, on skeletal points, ∇D_X is not a function in the usual sense but it can be regarded as a "multivalued function;" that is, if $P \in \mathrm{int}(X)$, then

$$\nabla D_X(P) = \left\{ \mathbf{v} \ : \ \mathbf{v} = \frac{P - Q}{\mathrm{d}(Q, P)}, \ Q \in P_C \right\}. \tag{2.3}$$

[2]See Calabi and Hartnett [11] for a higher level discussion without proofs.

Corollary 2.10 guarantees Equation 2.3 to give a function whenever $P \notin \mathrm{Sk}(X)$, but on skeletal points it does definitely not. Indeed, $\nabla D_X(P)$ for $P \in \mathrm{Sk}(X)$ is thus defined by Equation 2.3 as the collection of all limit points of converging sequences $\nabla D_X(P_n)$ where $P_n \notin \mathrm{Sk}(X)$ for all n and $P_n \to P$.

The previous remark shows that the distance function is not differentiable at skeletal points according to Definition 2.13. The same is true for almost all points of the medial axis as given by Definition 2.12; Theorem 2.14 makes this claim precise.

Theorem 2.14. *Let*

$$E = \{P \in X \ : \ P \textit{ is a centre of curvature for } \partial X \textit{and } |P_C| = 1\}.$$

Then, $\mathbf{MA}(X) - E = \mathrm{Sk}(X)$.

In other words, the skeleton set $\mathrm{Sk}(X)$ is the same as the medial axis set $\mathbf{MA}(X)$ except possibly for some end-points of $\mathbf{MA}(X)$. This is the intuitive interpretation assuming that the skeleton is just a collection of curves, i.e. assuming that it is thin. This assumption does hold but it is not a new result. Several authors, including Matheron [32] and Calabi (see See Calabi and Hartnett [11]), have analyzed the skeleton's thickness and equivalent formulations of the structure, but their frameworks are unnecessarily complicated for our purposes here. Therefore, rather than follow an approach based on those investigations, we shall instead provide self-contained proofs of both Theorem 2.14 and the following corollary (Corollary 2.15) in Appendix A; refer to Theorem A.2 and Corollary A.3.

Corollary 2.15. *Let X be a shape. The skeleton of X is a collection of bounded curves, i.e.*

$$\mathrm{int}(\mathrm{Sk}(X)) = \varnothing.$$

Remark 2.3. This result is not unique to planar shapes; in fact, it can be shown that the skeleton of an n-dimensional bounded manifold (defined by, say, maximal balls exactly similarly to Definition 2.12) is an $(n-1)$-dimensional structure (see Bouix [8] for $n = 3$).

The skeleton is not only thin, it is also topologically equivalent to the original shape. Specifically, there is a natural way to map holes in $\mathrm{Sk}(X)$ to holes in X in a one-to-one and onto fashion. In fact, a stronger relationship exists: there is a continuous process that takes the boundary of the shape to its skeleton. This process is given formally by a *homotopy* function.

Definition 2.16. Let $\gamma_1 : \mathbb{R} \to \mathbb{R}^2$ and $\gamma_2 : \mathbb{R} \to \mathbb{R}^2$ be two curves. These curves are *homotopic* if and only if there is a *homotopy*[3] $H : [0,1] \times \mathbb{R} \to \mathbb{R}^2$ which is a continuous function and such that $H(0,t) = \gamma_1(t)$ and $H(1,t) = \gamma_2(t)$.

Thus, it can be shown that the shape's boundary (actually much less restrictive shapes as those assumed in this thesis) is homotopic to the skeleton (refer to Lieutier [29] or Calabi and Hartnett [11]). This is the meaning of the following:[4]

Theorem 2.17. *Let X be a shape and* $\mathrm{Sk}(X)$ *its skeleton. Then X and* $\mathrm{Sk}(X)$ *are of the same homotopy type.*

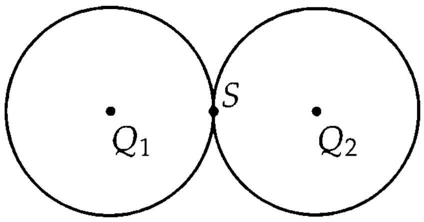

FIGURE 2.2. Necessity of condition (i) in Definition 2.6. The circles meet at a single point S, so that their interiors are disconnected. The skeleton of the shape is then the two points Q_1 and Q_2. Thus, even though the shape is connected, its skeleton is not and, therefore, it is not of the same homotopy type.

Remark 2.4. This result holds under the assumptions for shapes as specified by Definition 2.6 and, as Figure 2.2 illustrates, justifies condition (i)—if the interior of the shape is not connected, then the theorem need not be true. In fact, the relationship is somewhat stronger: Lieutier [29] gives a proof that the skeleton of any bounded open subset of $\mathbb{R}^n$ is homotopy equivalent to its skeleton. Thus, if the shape is not of the same homotopy type as its interior, it will not be homotopic to its skeleton (that is, $\mathrm{Sk}(X)$).

Theorem 2.17 suggests that boundary points are mapped to skeleton points, but what can be said about mapping skeleton points to boundary points? It turns out that the process is almost invertible—it fails only for sharp convex corners of the boundary. Proposition 2.19 formalizes the result (see Proposition A.4 in Appendix A for the proof).

[3] A more general definition may be found in a standard book on Topology; e.g. see Munkres [40].
[4] See footnote 3.

Résumé

L'axe médian de formes 2D est l'ensemble de points où la fonction Euclideinne de distance au contour n'est pas differentiable. Cette structure, aussi connue par le nom de squelette, est une très bonne representation: elle est mince, homotopique à la forme, invariante aux transformations rigides et elle a une structure combinatoire naturelle. Cependant, malgré tous ces benefices, le développement d'algorithmes qui calculent le squelette d'une forme 2D à été dificile.

Dans cette thèse, nous présentons une solution en deux étapes. Premièrement, nous montrons que la limite du flux moyen à travers une region tendant vers un point sur l'axe médian est differente que celle si le point n'appartient pas à cette structure. Nous obtenons ainsi un critère pour identifier les points du squelette. Deuxièmement, en étudiant le comportement du flux moyen dans le cas discret, nous adaptons ce critère de sélection et développons une procédure qui approxime l'axe médian d'une forme dont l'erreur est arbitrairement petite.

Definition 2.18. Let X be shape and $Q \in \partial X$ be a point where the boundary is C^0 but not C^1. Let θ be the angle inside X formed by the tangents to ∂X taken while approaching Q from either side. If $\theta < \pi$ then Q is a *sharp corner* of the boundary; see Figure 2.3. If $\theta > \pi$ then Q is a *dull corner*.

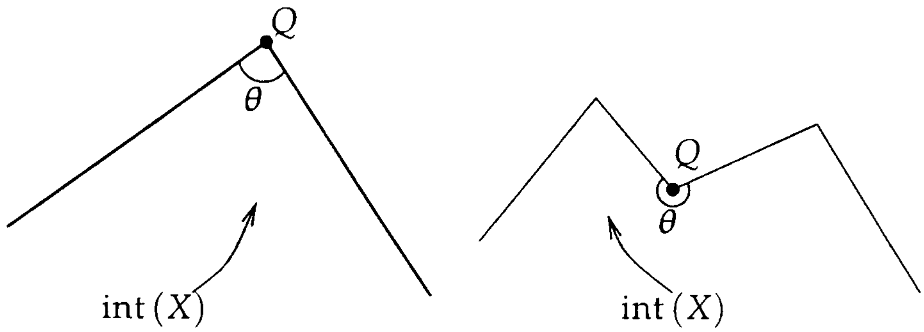

FIGURE 2.3. An examples of a sharp corner (LEFT), and a dull corner (RIGHT); see Definition 2.18.

Proposition 2.19. *Let X be a shape and $Q \in \partial X$. Then, there exists a $P \in \mathrm{Sk}(X)$ such that $Q \in P_C$ or Q is a sharp corner of the boundary.*

So now we know something about relating skeletal points to boundary points, but what we are really after is a way to make the skeleton into a graph. To do this, let us classify the points in $\mathrm{Sk}(X)$ according to their boundary support.

Definition 2.20 (Skeletal Points). Let X be a shape and $P \in \mathrm{Sk}(X)$. Refer to Figure 2.4 for examples.

i. P is called a *regular point* if $|P_C| = 2$.

ii. P is an *end-point* if there exists $\varepsilon_0 > 0$ such that for any $0 < \varepsilon < \varepsilon_0$ the circle centered at P intersects $\mathrm{Sk}(X)$ at a single point.

iii. P is a *junction point* if P_C has three or more connected components.

iv. P is a *pseudo-junction point* if P_C contains infinitely many points but only two connected components.

Theorem 2.21. *The skeleton of a shape X is well behaved:*

(1) Most skeletal points are regular; that is, there are only finitely many skeletal points which are not regular.

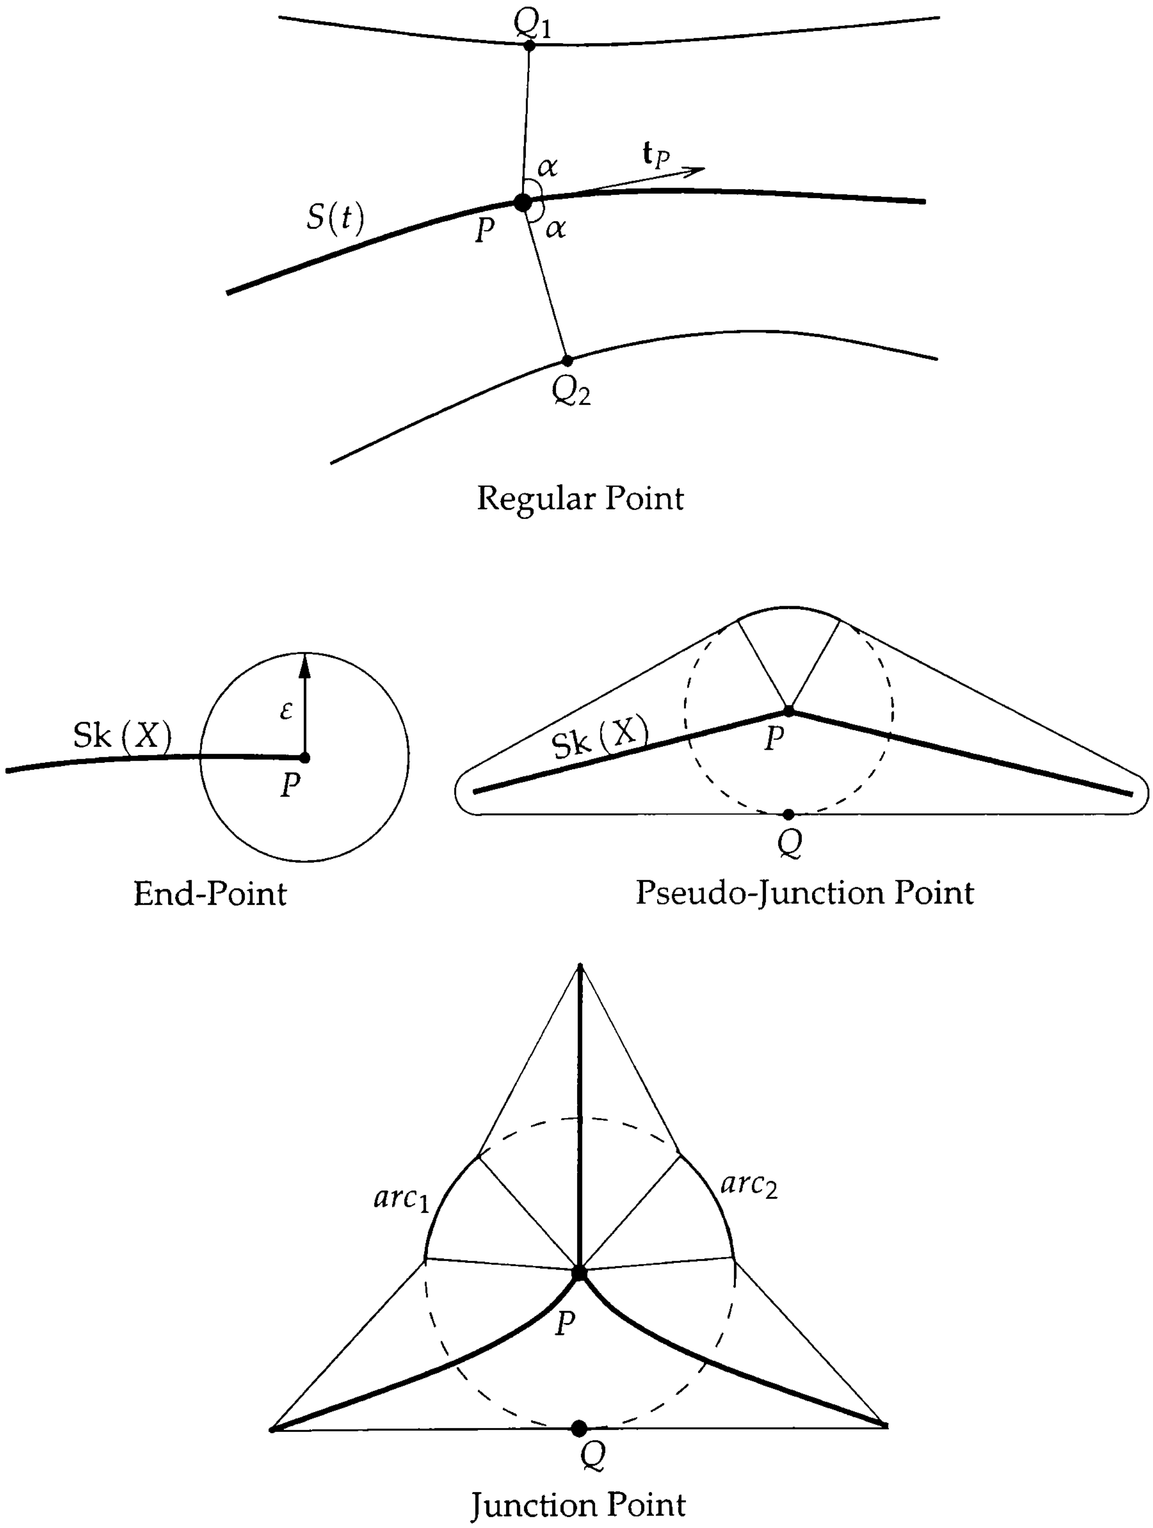

FIGURE 2.4. Examples of the different kinds of skeletal points; see Definition 2.20. The isolated points on the boundary (denoted Q, Q_1, Q_2) together with the bolder arcs on the boundary form the boundary support (denoted P_C) for the skeletal point P.

(2) *The skeleton is*

$$\mathrm{Sk}(X) = \bigcup_{i=1}^{n} \mathrm{Image}(S_i(t))$$

where $S_i(t) : [0,1] \to \mathrm{Sk}(X)$ is a curve of regular points and of maximum length. Each such curve is of finite length and n is a natural number. Each $S_i(t)$ is C^1 everywhere and C^k almost everywhere; there are only finitely many points where the curve is not C^k. Further, $S_i(t)$ do not self-intersect and $S_i(t)$ and $S_j(t)$ may meet only at their extreme points, whenever $i \neq j$.

(3) *If P is a regular skeletal point and its boundary support P_C contains a dull corner, then the unique skeletal curve $S(t)$ passing through it is strictly C^1 (i.e. it is not C^2) if and only if any neighborhood of P on $S(t)$ contains a point with boundary support that does not have any dull corners.*

(4) *If $r_i(t)$ is defined such that $(S_i(t), r_i(t)) \in \mathbf{ST}(X)$ and $S_i(t)$ as in Part 2, then $r_i(t)$ is as differentiable as $S(t)$, i.e. it is C^k where $S_i(t)$ is and it is strictly C^1 where $S_i(t)$ is.*

(5) $\mathrm{Sk}(X)$ *is path connected.*

(6) *If $P \in \mathrm{Sk}(X)$ is a junction point, then P_C has n connected components for some integer $n > 2$. Further, exactly n distinct skeletal curves meet at P.*

This result is a consequence of our development so far and several theorems in Choi *et al.* [13].

Definition 2.22. The curves $S(t)$ as in Theorem 2.21(2) are called *skeletal curves* and the $r(t)$ corresponding to an $S(t)$ is the *radius function* along the skeletal curve.

Leyton [28] demonstrated how each skeletal curve is related to the boundary. He called this result the Symmetry-Curvature Duality Theorem which can be interpreted as the following

Theorem 2.23 (Symmetry-Curvature Duality). *Let X be a shape. To each skeletal curve there correspond exactly two subcurves of the boundary ∂X.*

Thus, (finally) we can realize the skeleton as a graph. To do this, notice that there are only finitely many skeletal curves in $\mathrm{Sk}(X)$, they do not self-intersect and are mutually disjoint except at their extremities. The combinatorial structure can then be defined by taking the end-points and junction points as nodes, and the curves connecting these, as the edges. Theorem 2.17 implies that, for connected shapes, the graph will be connected as well. Further, if the shape has no holes, the graph must be a tree.

However, the most important property of this graph, from a computational standpoint, is that it is *finite*. The fact that the boundary of a shape must be given by finitely many real analytic curves (see Definition 2.6(ii)) is the reason for this behavior. Matheron [32] gives a counterexample of a shape with boundary made of infinitely (but countably) many real analytic curves where there are infinitely many skeletal curves. On the other hand, there are shapes with only two smooth (i.e. C^∞) curves defining their boundary but which are not real analytic such that the skeleton is an infinite graph; see Choi *et al.* [13].

In summary, we have seen that the skeleton of shapes under our assumptions must be thin and with finite combinatorial structure. Therefore, it should be possible to implement the medial axis for such objects on a digital machine, but exactly how much information does this thin structure carry about the original shape? In the following section, we study the adequacy of the skeleton as a descriptor for shape.

2.4 Boundary Representation through ST (X)

In this section, we shall demonstrate that the skeleton transform **ST** (X) of a shape X is a faithful representation, equivalent to the original object. Our analysis will reveal that the curvature of the boundary can be obtained from the **ST** (X) and how this structure can be conceived of as an invariant of the group of rigid transformations in the plane. We shall also introduce the notion of α-skeletons which are restrictions of **ST** (X).

We begin by studying the second component of the skeleton transform, the radius function associated to a skeletal curve. Denote the pair by $(S(t), r(t)) \in$ **ST** (X) and observe that the $r(t)$ is simply a restriction of the distance function D_X to the points on $S(t)$ with the sign is ignored. Thus, the derivative of $r(t)$ should be somehow related to the gradient of D_X. Indeed, although not a function on regular points of the skeleton, ∇D_X can be seen as a pair of vectors on such locations and $r'(t)$ captures their geometric relationship. Figure 2.5 illustrates this and Theorem 2.24 (with proof in Appendix A, Theorem A.5) formalizes the claim.

Theorem 2.24. *Let P be a regular skeletal point and denote by Q_1 and Q_2 the two distinct points in P_C. Let $S(t)$ be the skeletal curve that that passes through P, such that $S(t_0) = P$. Then,*

(a) the angle $\angle Q_1PQ_2$ is bisected by $S'(t_0)$, i.e. the tangent to $S(t)$ at t_0; and

*(b) $|r'(t_0)| = \cos\alpha$, where $r(t_0)$ is the radius function at P (i.e. $P = S(t_0)$ and $(S(t_0), r(t_0)) \in$ **ST** (X)) and $\alpha \leq \frac{\pi}{2}$ is half of $\angle Q_1PQ_2$.*

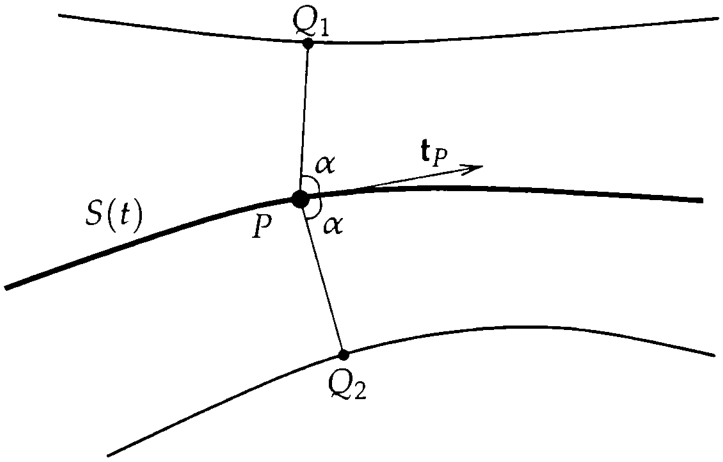

FIGURE 2.5. The object angle $\alpha = \alpha(P)$ at a regular skeletal point P. Here $S(t)$ is a parameterization of the skeleton curve that passes through P, i.e. $P = S(t_0)$, and $\mathbf{t}_P = S'(t_0)$ is the tangent at t_0. See Theorem 2.24.

Recall that skeletal curves consist of regular points except for their two extremities (see Theorem 2.21(2)). Therefore, the angle α in Theorem 2.24 is defined on the whole curve and can be thought of as a function.

Definition 2.25. Let $S(t)$ be a skeletal curve and let $\alpha(t)$ be as in Theorem 2.24. The function $\alpha(t)$ is called the *object angle* function.

Thus, the object angle at a regular point is the angle formed by the tangent to the skeleton in the direction of decreasing radius function and the gradient vector on one side of the curve. In particular, the skeleton point P, the radius value at P and the object angle there are sufficient to recover the two points on the boundary which are closest to P; that is, using the notation on Figure 2.5, we can write

$$Q_i = P + |D_X(P)|\, R((-1)^{i+1}\alpha)\mathbf{t}_P, \qquad i = 1, 2 \tag{2.4}$$

where $R((-1)^{i+1}\alpha)$ represents counterclockwise rotation by plus α to obtain Q_1, and by minus the object angle to recover Q_2.

Now recall that there is a continuous process that maps boundary points to skeleton points (see Theorem 2.17). Thus, if P_1 and P_2 are regular points on the same skeletal curve and are close, then their boundary support should also be close. In practice, this means that Equation 2.4 defines the Q_is continuously along $S(t)$ and, consequently, it regenerates the two contour curves (whose existence is guaranteed by Theorem 2.23) corresponding to $S(t)$. Formally, this translates into the following theorem.[5]

Theorem 2.26 (Reconstruction). *Let X be a shape. Let $S(t)$ be a skeletal curve and denote by $C_i : \mathbb{R} \to \partial X$ the two curves corresponding to it. Then, the boundary curves may be parameterized given a parameterization of $S(t)$. In particular,*

$$C_i(t) = S(t) + r(t)\left(-r'(t)S'(t) + (-1)^{i+1}\sqrt{1-(r'(t))^2}S'_{\perp}(t)\right), \qquad i = 1,2 \tag{2.5}$$

where $S'(t)$ is the unit tangent to the skeleton and $S'_{\perp}(t)$ is obtained by a counterclockwise rotation of $S'(t)$ by $\pi/2$.

The result is significant because it provides an explicit relationship between the boundary and the skeleton. By the definition of Sk (X) (see Definition 2.13) it is clear that each skeletal point corresponds to at least two boundary points. Proposition 2.19, on the other hand, guarantees that all boundary points except sharp corners (of which there are only finitely many) correspond to at least one skeletal point. Therefore, Theorem 2.26 shows how to recover all points on the boundary except the sharp corners and demonstrates the efficiency of the skeleton transform as a representation for shape. However, it would not be an adequate representation if rigid transformations of the shape could change it drastically.

Fortunately, the **ST** (X) remains invariant under rigid transformation of X. This property may be expressed as the following theorem

Theorem 2.27 (Invariance). *Let τ be a rigid transformation of the plane $\mathbb{R}^2$ and X be a shape. Then, $\tau\mathrm{Sk}(X) = \mathrm{Sk}(\tau X)$.*

In other words, the rigid transformations commute with the skeletonization operator Sk $(\cdot)$. The claim is a direct consequence of the Reconstruction Theorem:

[5] A formal proof is given by Choi *et al.* [13].

a regular skeletal point P and its two corresponding boundary points Q_1, Q_2 define a triangle which is invariant under rigid transformations, in particular, the relative positions of the gradient lines emanating from Q_1 and Q_2 and meeting at P (i.e. the sides of the isosceles triangle) remain unchanged.

Another consequence of the Reconstruction Theorem is the fact that boundary curvature may be expressed by second order differential properties of the skeleton transform. We have the following theorem (shown as Theorem A.6 in Appendix A).

Theorem 2.28. *Let $S(t)$ be a segment of a skeletal curve such that $r(t)$ is monotonically decreasing and $\|S'(t)\| = 1$. Assume further that $S(t)$ is at least C^2 and let $C(t)$ be a contour reconstruction according to Equation 2.5. Then, the curvature of the boundary segment $C(t)$, denoted $\kappa_C(t)$, is*

$$\begin{aligned} |\kappa_C(t)| &= \frac{|\alpha'(t) - \kappa_S(t)|}{\|C'(t)\|} \\ &= \left| \frac{\alpha'(t) - \kappa_S(t)}{r(t)(\alpha'(t) + \kappa_S(t)) - \sin\alpha(t)} \right| \end{aligned}$$

It turns out that the quantity in the denominator actually comes up from a well known boundary-skeleton correspondence—the boundary-to-axis ratio (see Blum and Nagel [7]).

Definition 2.29. The *boundary-to-axis ratio* is the ratio of a length on the boundary to the length of the corresponding skeletal portions. The notion makes sense, almost everywhere, locally as well.

This ratio is an indication as to how important a particular skeletal points is, that is, "how much" of the boundary it represents. Indeed, not all medial points are perceptually relevant to the shape, perhaps even whole branches. Another interesting way of assigning relevance is through the amount of area represented. Define the following simplification of a skeleton:

Definition 2.30. The *α-skeleton*, denoted Sk (X, α), of a shape X is a restriction of the skeleton of X, i.e. Sk (X), obtained by retracting the end-points of Sk (X) until a skeletal point is reached with object angle greater than or equal to α. Alternatively, letting Γ_P be the set of shortest paths starting at P and ending at an end-point of

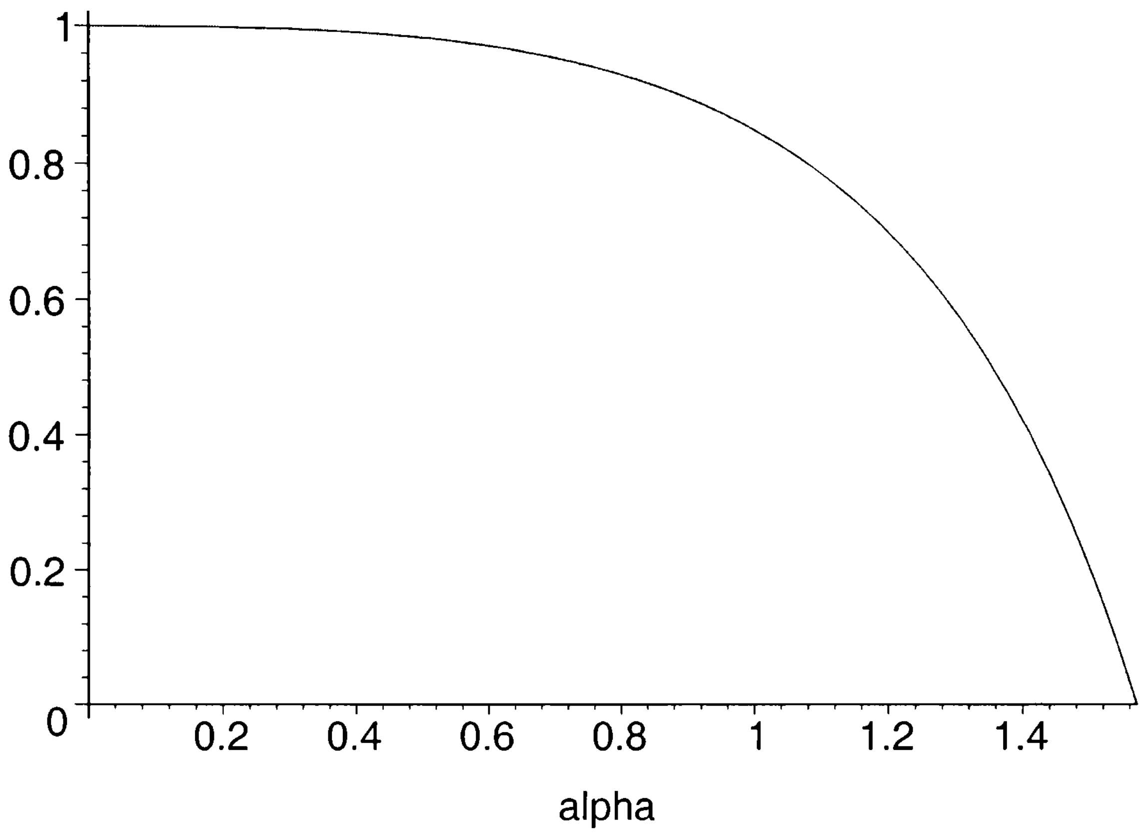

FIGURE 2.6. The proportion of area represented by an α-skeleton. Here, $\alpha \in [0, \pi/2]$.

Sk (X), the α-skeleton is formally given by

$$\mathrm{Sk}\,(X, \alpha) = \left\{P \in \mathrm{Sk}\,(X) \ : \ \forall \gamma \in \Gamma_P, \exists Q \in \gamma \text{ such that } \alpha_Q \geq \alpha\right\}$$

where α_Q denotes the object angle at point Q.

The *α-skeleton transform*, denoted **ST** (X, α), is the restriction of **ST** (X) induced by Sk (X, α).

Theorem 2.31 (α-Skeleton). *Let $A(X)$ denote the area of shape X and $A(X, \alpha)$ the area of the shape represented by the α-skeleton transform of X, then*

$$\frac{A(X,\alpha)}{A(X)} \leq \frac{\pi}{\pi + \tan(\alpha) - \alpha} \tag{2.6}$$

Figure 2.6 is a plot of this bound. Notice that most of the area of the shape is preserved by an α-skeleton for α as high as 57° (i.e. 1rad). Thus, many of the skeletal branches of the medial axis of a polygonal approximation may be removed without affecting the represented shape much.

2.5 Summary

Before anything can be done about computing the skeleton of a shape, we must make precise exactly what the problem is. In this chapter, we discussed a mathematical definition for shape and formally introduced the concept of skeletonization. We saw that small variations in the assumptions about shape may have a big impact on the behavior of the skeleton. In particular, we saw that, under Definition 2.6, Sk (X) is a thin structure and it is readily interpreted as a finite graph. Furthermore, the skeleton transform was demonstrated to be a very good representation for shape, invariant to rigid transformations and providing a parameterization of the boundary. Finally, we saw that it is possible to simplify the skeleton by defining the α-skeleton which represents the original shape quite closely for small enough α.

Chapter 3

The AOF Criterion

Now that the formal introduction of shapes and their skeletons is in place, we can turn our attention to the task of identifying skeletal points. We shall proceed in two steps: first, we shall work in the continuous domain (i.e. all of $\mathbb{R}^2$) and, then, we shall specialize to the discrete lattice. In this chapter we shall assume the former setting and derive a criterion to classify points in a shape into either medial (on the skeleton) or non-medial. Even though the results will suggest a discrete algorithm, there are several issues which must be addressed in order to guarantee accuracy and robustness. Therefore, we shall postpone the presentation of such a procedure until Chapter 4.

The skeletal point identification criterion will be derived by studying the behavior of the distance function gradient field induced by the shape. We shall begin, in Section 3.1, by reviewing some basic definitions and the Divergence Theorem. Then, in Section 3.2, we shall extend this theorem—so it can be applied in the presence of skeletal points—and the criterion will be derived in Section 3.3 as a consequence of this extension. We shall then see (in Section 3.4) exactly how much information this criterion can give about skeletal points.[1]

3.1 Preliminaries

The results presented in this chapter will follow from a nonstandard application of the Divergence Theorem but, before going any further, we need to introduce some basic concepts and notations.

[1]The results presented in this chapter are largely those reported in Dimitrov *et al.* [17].

Contribution of Authors

Chapter 3 contains results published in collaboration with James Damon and Kaleem Siddiqi; see Dimitrov *et al.* [17]. The author's contributions include the derivation for circular regions (Section 3) and the experiments section.

Chapter 4 is an outgrowth of previously published work in collaboration with Carlos Phillips and Kaleem Siddiqi; see Dimitrov *et al.* [18].

Definition 3.1. The *outward flux* of a vector field $\mathbf{F}$ through a closed curve $\gamma \subset \mathbb{R}^2$ is given as the line integral

$$\mathcal{F}_\gamma = \int_\gamma \langle\, \mathbf{F}, \mathcal{N} \,\rangle \, ds$$

where $\mathcal{N}$ denotes the outward normals to γ and s is some parameterization of the curve. In the case of $\mathbf{F} = \nabla D_X$ and $\gamma = C_\varepsilon^P$ (the circle of radius ε centered at P), we shall denote the outward flux by $\mathcal{F}_\varepsilon(P)$.

Definition 3.2. The *average outward flux* of a vector field $\mathbf{F}$ through a closed curve γ is the the outward flux normalized by the length of the curve,i.e.

$$\mathscr{F}_\gamma = \frac{\mathcal{F}_\gamma}{L(\gamma)} = \frac{\int_\gamma \langle\, \mathbf{F}, \mathcal{N} \,\rangle \, ds}{\int_\gamma ds}.$$

Similarly to Definition 3.1, whenever $\mathbf{F} = \nabla D_X$ and $\gamma = C_\varepsilon^P$ we shall denote the average outward flux $\mathscr{F}_\varepsilon(P)$.

Definition 3.3. The divergence of a vector field $\mathbf{F} = \nabla D_X$, denoted $\operatorname{div}(\mathbf{F})$ or $\nabla \cdot \mathbf{F}$, is

$$\operatorname{div}(\mathbf{F}) = \frac{\partial^2 D_X}{\partial x^2} + \frac{\partial^2 D_X}{\partial y^2}.$$

Theorem 3.4 (Divergence Theorem). *Let X be a shape and let R be a path-connected region where $\operatorname{div}(\mathbf{F})$ is defined. Then, total divergence over R of $\mathbf{F}$ is equal to the the outward flux through the boundary of the region ∂R, i.e.*

$$\int_R \operatorname{div}(\mathbf{F}) \, dv \equiv \int_{\partial R} \langle\, \mathbf{F}, \mathcal{N} \,\rangle \, ds. \tag{3.1}$$

where dv is an area element.

Remark 3.1. It should be noted that this formulation of the Divergence Theorem is far from the general result—it is a specialization for regions in $\mathbb{R}^2$. A discussion of the n-dimensional case can be found in any standard text on smooth manifolds such as Warner [55].

3.2 An Extension of the Divergence Theorem

We shall now develop an extension of the Divergence Theorem which can be applied to investigate properties of the vector field $\mathbf{F} = \nabla D$ at skeletal points, where

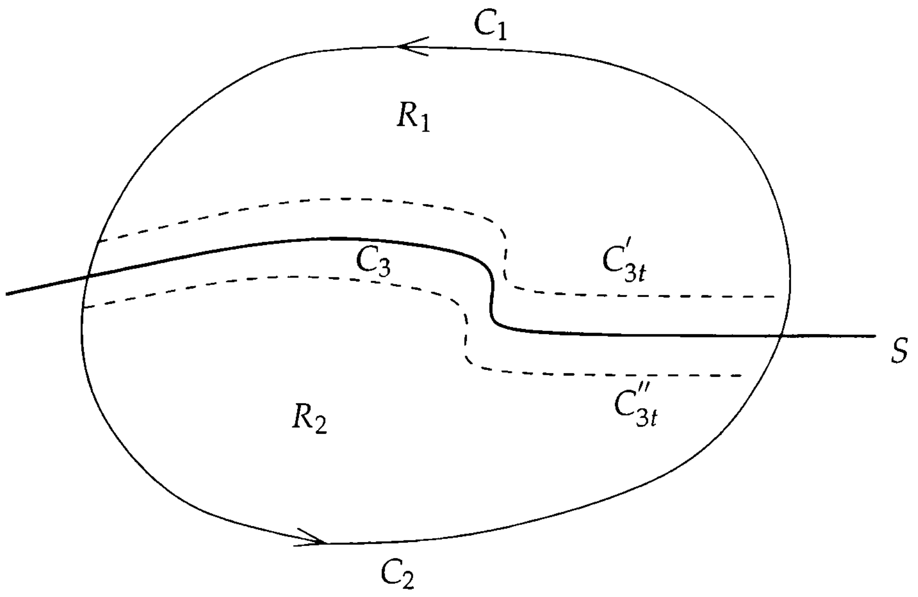

FIGURE 3.1. A region R which intersects a branch of the skeleton S.

it is discontinuous. Figure 3.1 illustrates the set up for the calculation which follows.

Let S be a branch of the skeleton and let $R = R_1 \cup R_2$ be a path connected region which intersects it. Let $\partial R = C_1 \cup C_2$ and $C_3 = S \cap R$. Let C'_{3t}, C''_{3t} be parallel curves to C_3 which approach C_3 as $t \to 0$. Let R_{1t} and R_{2t} be the regions obtained from R_1 and R_2 by removing the region between the curves C'_{3t} and C''_{3t} Finally, let $\mathbf{F}_+$ denote $\mathbf{F}$ above S and $\mathbf{F}_-$ denote $\mathbf{F}$ below S.

The outward flux of $\mathbf{F}$ through ∂R is given by

$$\int_{\partial R} \langle\, \mathbf{F}, \mathcal{N} \,\rangle \; ds = \int_{C_1} \langle\, \mathbf{F}, \mathcal{N} \,\rangle \; ds + \int_{C_2} \langle\, \mathbf{F}, \mathcal{N} \,\rangle \; ds.$$

Applying the Divergence Theorem to R_{1t} and R_{2t}

$$\int_{R_{1t}} \operatorname{div}(\mathbf{F}) \; dv = \int_{C_{1t}} \langle\, \mathbf{F}, \mathcal{N} \,\rangle \; ds + \int_{C'_{3t}} \langle\, \mathbf{F}, \mathcal{N} \,\rangle \; ds,$$

$$\int_{R_{2t}} \operatorname{div}(\mathbf{F}) \; dv = \int_{C_{2t}} \langle\, \mathbf{F}, \mathcal{N} \,\rangle \; ds + \int_{-C''_{3t}} \langle\, \mathbf{F}, \mathcal{N} \,\rangle \; ds.$$

Adding the above two equations we have

$$\begin{aligned}\int_{R_{1t}} \operatorname{div}(\mathbf{F})\, dv + \int_{R_{2t}} \operatorname{div}(\mathbf{F})\, dv &= \int_{C_{1t}} \langle\, \mathbf{F}, \mathcal{N}\,\rangle\ ds + \int_{C_{2t}} \langle\, \mathbf{F}, \mathcal{N}\,\rangle\ ds \\ &+ \int_{C'_{3t}} \langle\, \mathbf{F}, \mathcal{N}\,\rangle\ ds + \int_{-C''_{3t}} \langle\, \mathbf{F}, \mathcal{N}\,\rangle\ ds.\end{aligned}$$

By Theorem 2.24, the tangent to the skeleton bisects the the angle between $\mathbf{F}_+$ and $\mathbf{F}_-$ at a regular skeletal point (see Figure 3.2). Thus, on C_3 we have

$$\langle\, \mathbf{F}_+, \mathcal{N}_+\,\rangle = \langle\, \mathbf{F}_-, \mathcal{N}_-\,\rangle, \tag{3.2}$$

where $\mathcal{N}_+, \mathcal{N}_-$ denote the normals to C_3 from above and from below, respectively. Thus, one can take the limit as $t \longrightarrow 0$ of both sides of the above equation to obtain the following extension of the Divergence Theorem

Theorem 3.5. *For a path connected region R which contains part of a skeletal curve S, the divergence of the vector field $\mathbf{F}$ is related to its flux through ∂R by the following equation*

$$\int_{R=R_1\cup R_2} \operatorname{div}(\mathbf{F})\, dv = \int_{\partial R} \langle\, \mathbf{F}, \mathcal{N}\,\rangle\ ds + 2\int_{C_3} \langle\, \mathbf{F}, \mathcal{N}_{C_3}\,\rangle\ ds.$$

Although the Divergence Theorem fails for such regions because $\mathbf{F}$ is discontinuous on C_3, the last integral is well defined due to Equation 3.2; we either take $\mathbf{F} = \mathbf{F}_+$ and $\mathcal{N}_{C_3} = \mathcal{N}_+$ or $\mathbf{F} = \mathbf{F}_-$ and $\mathcal{N}_{C_3} = \mathcal{N}_-$. Because $\langle\, \mathbf{F}, \mathcal{N}_{C_3}\,\rangle$ is in the interval $(0, 1]$ on the skeleton, it also follows that

$$\int_{\partial R} \langle\, \mathbf{F}, \mathcal{N}\,\rangle\ ds \geq \int_{R} \operatorname{div}(\mathbf{F})\, dv - 2L(C_3).$$

where $L(C_3)$ denotes the Euclidean length of the curve.

3.3 Skeletal Point Selection Criterion

We now consider the limit values of the outward flux and the average outward flux, of the vector field $\mathbf{F} = \nabla D$ through a convex curve that is the boundary ∂R of a region R, as the region shrinks to a point. The results reported here are actually a special case of a more general result which applies to the case of any shrinking convex domain of arbitrary (but finite) dimension; see Damon [14].

We begin by considering the case where the limit point P does not lie on the skeleton and hence Equation 3.1 applies. We can write

$$\int_R \operatorname{div}(\mathbf{F})\, dv \equiv \int_0^L \langle\, \mathbf{F}(P) + \delta(s),\, \mathcal{N} \,\rangle\, ds$$

where L is the Euclidean length of ∂R, $\mathbf{F}(P)$ is the value of the vector field at the limit point P and $\delta(s)$ is the quantity added to get the value of $\mathbf{F}$ at neighboring points on ∂R ($\mathbf{F}$ is continuous at P). In the limit as $L \to 0$ we have

$$\lim_{R \to P} \mathcal{F}_{\partial R} = \lim_{L \to 0} \left(\int_0^L \langle\, \mathbf{F}(P),\, \mathcal{N} \,\rangle\, ds + \int_0^L \langle\, \delta(s),\, \mathcal{N} \,\rangle\, ds \right). \tag{3.3}$$

Now, $\langle\, \mathbf{F}(P),\, \mathcal{N} \,\rangle$ is a continuous function over the ∂R for a small enough R. The boundary ∂R is also a closed curve, so, by the Fundamental Theorem of Calculus, the first integral is identically equal to zero for all non-skeletal points. The limit of the second integral must also be zero because, as Corollary 2.10 guarantees, $\delta(s)$ must be continuous and the Fundamental Theorem of Calculus applies again. The average outward flux is shown to vanish at non-medial points exactly similarly—one must divide by the length of ∂R which also tends to zero but the limit is determined by the fact that the numerator is identically zero.

We now consider the second case where the limit point is a skeletal point and hence Theorem 3.5 applies, which we rewrite as

$$\mathcal{F}_{\partial R} = \int_R \operatorname{div}(\mathbf{F})\, dv - 2 \int_{C_3} \langle\, \mathbf{F},\, \mathcal{N}_{C_3} \,\rangle\, ds,$$

which implies that the average outward flux can be seen as

$$\mathscr{F}_{\partial R} = \frac{\int_R \operatorname{div}(\mathbf{F})\, dv}{L(\partial R)} - \frac{2 \int_{C_3} \langle\, \mathbf{F},\, \mathcal{N}_{C_3} \,\rangle\, ds}{L(\partial R)}$$

Considering the limit as R shrinks to a point, the argument for non-skeletal point applies to the first term on the right hand side. Thus, the potentially nonzero term is the second one. Therefore,

$$\lim_{\partial R \to P} \mathscr{F}_{\partial R} = -\lim_{\partial R \to P} \frac{2 \int_{C_3} \langle\, \mathbf{F},\, \mathcal{N}_{C_3} \,\rangle\, ds}{L(\partial R)} = -\lim_{\partial R \to P} \frac{2 \int_{C_3} \langle\, \mathbf{F}(P) + \delta(s),\, \mathcal{N}_{C_3} \,\rangle\, ds}{L(\partial R)}$$

where P is the limit point and $\delta(s)$ is the quantity added to get the value of $\mathbf{F}$ at neighboring points on C_3. Owing to the fact that the integrand is in $(0,1]$, the value of this integral is bounded as

$$\lim_{\partial R \to P} \frac{-2\left(\sup_{C_3} \langle \mathbf{F}_{C_3}, \mathcal{N}_{C_3} \rangle\right) L(C_3)}{L(\partial R)} \leq \lim_{\partial R \to P} \mathscr{F}_{\partial R}$$

$$\lim_{\partial R \to P} \mathscr{F}_{\partial R} \leq \lim_{\partial R \to P} \frac{-2\left(\inf_{C_3} \langle \mathbf{F}_{C_3}, \mathcal{N}_{C_3} \rangle\right) L(C_3)}{L(\partial R)}$$

Thus the average outward flux through a region shrinking to a regular skeletal point is

$$\lim_{\partial R \to P} \mathscr{F}_{\partial R} = -2 \langle \mathbf{F}(P), \mathcal{N}_{C_3}(P) \rangle \lim_{R \to P} \frac{L(C_3)}{L(\partial R)} \,. \tag{3.4}$$

Summarizing the above results, we have the property that whereas the limit value of the outward flux is zero for both skeletal[2] and non-skeletal points, the average outward flux has a different limiting behavior at skeletal points than at non-skeletal ones, providing a theoretical justification for its use in the Hamilton-Jacobi skeletonization algorithm Siddiqi *et al.* [48].

3.4 Circular Neighborhoods

We now specialize the average outward flux calculation to the case of circular neighborhoods shrinking to a skeletal point. Instead of using Equation 3.4 to study the limiting behavior, we shall exploit the properties of this special case. Doing so will not only corroborate the analysis of Section 3.3, but it will allow us to study regular points as well as extreme points of skeletal curves.

We shall treat the three cases (see Definition 2.20) of regular points, junction points, and end-points of the skeleton separately.

3.4.1 Regular Skeletal Points

A *regular skeletal point* P is one for which $P_C = \{Q_1, Q_2\}$ for $Q_1 \neq Q_2$. Let $\mathbf{n}_1$ and $\mathbf{n}_2$ be the unit inward normals to the boundary at Q_1 and Q_2 respectively. Let $\mathbf{t}_P$ be the unit tangent vector to the skeleton at P and define the *object angle* at P to be $\alpha(P) \in [0, \pi/2]$, such that

$$\mathbf{n}_1 \cdot \mathbf{n}_2 = \cos 2\alpha(P).$$

[2]Multiply Equation 3.4 by $L(\partial R)$ to get $\lim_{R \to P} \mathcal{F}_{\partial R} = -2 \langle \mathbf{F}(P), \mathcal{N}_{C_3}(P) \rangle \lim_{R \to P} L(C_3) = 0$.

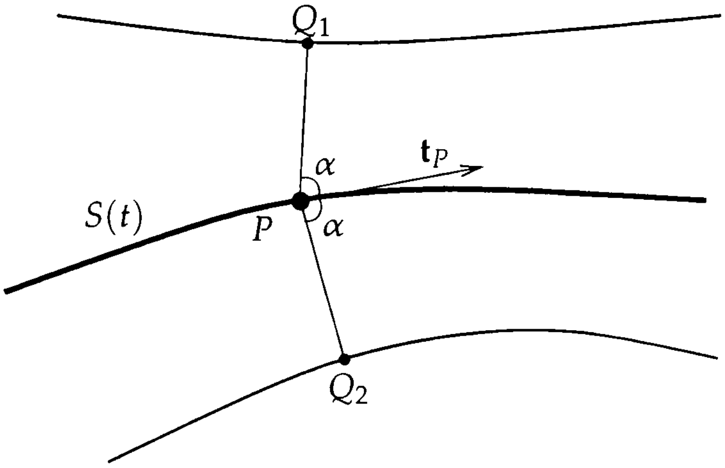

FIGURE 3.2. The object angle $\alpha = \alpha(P)$ at a regular skeletal point P. Here $S(t)$ is a parameterization of the skeleton curve. Hence, $\mathbf{t}_P = S'(t_0)$ is the tangent at t_0, i.e. where $P = S(t_0)$.

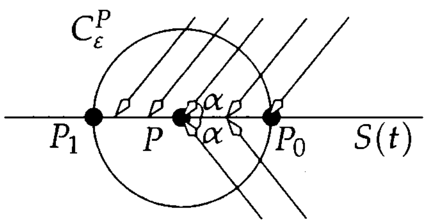

FIGURE 3.3. The distance function gradient vector field in the ε-neighborhood of P is given by a step function – one value for the "top" semi-circle and another for the "bottom" one. Both these vectors form an angle of $\alpha = \alpha(P)$ with $\mathbf{t}_P$, since the skeleton is assumed to cut C_ε^P in half at P_0 and P_1.

It follows that $\mathbf{n}_i \cdot \mathbf{t}_P = \cos\alpha(P)$ for $i = 1, 2$ (see Figure 3.2).

Now, let C_ε^P be the circle with radius ε centered at P. Let $C_\varepsilon^P : [0, 2\pi\varepsilon] \to \mathbb{R}^2$ be arc-length parameterized as

$$C_\varepsilon^P(s) = \varepsilon\left(\cos\left(\frac{s}{\varepsilon} + \theta(\mathbf{t}_P)\right), \sin\left(\frac{s}{\varepsilon} + \theta(\mathbf{t}_P)\right)\right) + P, \tag{3.5}$$

where $C_\varepsilon^P(0) = P + \varepsilon\mathbf{t}_P$ and $C_\varepsilon^P(\pi\varepsilon) = P - \varepsilon\mathbf{t}_P$. Now consider Figure 3.3. Here, it is assumed that the gradient field has one value along $C_\varepsilon^P(s)$ for $s \in (0, \pi\varepsilon)$ and another for $s \in (\pi\varepsilon, 2\pi\varepsilon)$. Also, both $C_\varepsilon^P(0) = P_0$ and $C_\varepsilon^P(\pi\varepsilon) = P_1$ are on the skeleton. Let the outward normal of this circle at s be $\mathcal{N}(s)$. Hence, the outward flux of ∇D though $C_\varepsilon^P(s)$ is

$$\begin{aligned}\mathcal{F}_\varepsilon(P) &= \int_0^{2\pi\varepsilon} \Big\langle \nabla D(C_\varepsilon^P(s)), \mathcal{N}(s) \Big\rangle \, ds \\ &= -\varepsilon \int_0^{\pi} \cos(\alpha - s)\, ds - \varepsilon \int_\pi^{2\pi} \cos(-\alpha - s)\, ds \\ &= -4\varepsilon \sin(\alpha)\end{aligned} \tag{3.6}$$

Notice that this calculation holds regardless of the orientation of $\mathbf{t}_P$. However, it makes very strict assumptions that do not hold in most situations. Fortunately, the general case is similar to this one.

There are only two differences: (1) $C_\varepsilon^P(0)$ and $C_\varepsilon^P(\pi\varepsilon)$ may not be on the skeleton, and (2) the distance function gradient field may take on more than two values along $C_\varepsilon^P(s)$ for $s \in [0, 2\pi\varepsilon]$. For small enough ε, the circle will intersect the skeleton at precisely two points, which we label $P_0 = C_\varepsilon^P(\delta_0\varepsilon)$ and $P_1 = C_\varepsilon^P((\pi + \delta_1)\varepsilon)$. Thus, the distance function gradient field is continuous on $C_\varepsilon^P(s)$ for $s \in I_0 = (\delta_0\varepsilon, (\pi + \delta_1)\varepsilon)$ and also for $s \in I_1 = ((\pi + \delta_1)\varepsilon, (2\pi - \delta_0)\varepsilon)$ [3]. However, it may take on more than one value in the intervals I_0 and I_1. Define $\beta_0(s)$ and $\beta_1(s)$ on I_0 and I_1 respectively, to account for such eventualities:

$$\mathbf{t}_P \cdot \theta(C_\varepsilon^P(s)) = \cos\left(\alpha(P) + \beta_0(s)\right) \quad , s \in I_0$$

$$\mathbf{t}_P \cdot \theta(C_\varepsilon^P(s)) = \cos\left(-\alpha(P) + \beta_1(s)\right), \; s \in I_1.$$

Therefore, the outward flux calculation for regular skeletal points becomes

$$\begin{aligned}\mathcal{F}_\varepsilon(P) &= \int_0^{2\pi\varepsilon} \Big\langle \nabla D(C_\varepsilon^P(s)), \mathcal{N}(s) \Big\rangle \, ds \\ &= -\varepsilon \int_{\delta_0}^{\pi+\delta_1} \cos(\alpha + \beta_0(s) - s)\, ds \\ &\quad -\varepsilon \int_{\pi+\delta_1}^{2\pi-\delta_0} \cos(-\alpha + \beta_1(s) - s)\, ds.\end{aligned}$$

The continuity of the distance function gradient field along the circle implies

[3] However, it is not necessarily continuous on the closure of $I_0 \cup I_1$.

that both $\beta_0(s)$ and $\beta_1(s)$ are continuous functions. Further, as $\varepsilon \to 0$, necessarily

$$\lim_{\varepsilon \to 0} \sup_{s \in [\delta_0, \pi+\delta_1]} |\beta_0(s)| = 0$$

$$\lim_{\varepsilon \to 0} \sup_{s \in [\pi+\delta_1, 2\pi-\delta_0]} |\beta_1(s)| = 0.$$

Also, since the skeletal curve has continuous tangents, we must have that $\lim_{\varepsilon \to 0} \delta_i = 0$ for $i = 0, 1$. Therefore the average outward flux through a shrinking circular region is given by

$$\begin{aligned}
\lim_{\varepsilon \to 0} \frac{\mathcal{F}_\varepsilon(P)}{2\pi\varepsilon} &= -\lim_{\varepsilon \to 0} \frac{\varepsilon}{2\pi\varepsilon} \int_{\delta_0}^{\pi+\delta_1} \cos(\alpha + \beta_0(s) - s)\, ds \\
&\quad - \lim_{\varepsilon \to 0} \frac{\varepsilon}{2\pi\varepsilon} \int_{\pi+\delta_1}^{2\pi-\delta_0} \cos(-\alpha + \beta_1(s) - s)\, ds \\
&= -\frac{1}{2\pi} \int_0^{\pi} \cos(\alpha - s)\, ds - \frac{1}{2\pi} \int_{\pi}^{2\pi} \cos(-\alpha - s)\, ds \\
&= -\frac{2}{\pi} \sin\alpha.
\end{aligned} \tag{3.7}$$

In summary, we have shown that, using the notation from Equation 3.4, C_3 tends to the diameter of the circle and $L(C_3)/L(\partial R) = 2\varepsilon/2\pi\varepsilon = 1/\pi$, which implies that the limit of the average outward flux is, essentially, the object angle at the regular point. Thus, Equation 3.4 is corroborated for the case of shrinking circles. However, this exercise was not just a verification, we have introduced the notation for our subsequent analysis of the other types of skeletal points.

3.4.2 Skeletal End-Points

Let P be a skeletal end-point. Let the point Q_ε be on the branch which is at distance ε from P. Choose ε small enough so that Q_ε is a regular skeletal point. Thus, the object angle is well defined for Q_ε. Now, let

$$\alpha_P = \lim_{\varepsilon \to 0} \alpha(Q_\varepsilon).$$

This limit makes sense, because the circle[4] C_ε^P intersects the skeleton at a single point and the object angle function is continuous along a skeletal branch.

Now consider Figure 3.4. Along the arc arc_{α_P} opposite to the skeleton curve, the distance function gradient field must coincide with the inner normals of the circle. This is because the end-point results from the collapse of a circular arc (possibly a point if $\alpha_P = 0$) on the boundary. On the rest of the circle, the distance function gradient field behaves as if P were a regular skeletal point. Thus,

$$\begin{aligned}
\mathcal{F}_\varepsilon(P) = &- \varepsilon \int_{-\alpha_P}^{\alpha_P} ds \\
&- \varepsilon \int_{\alpha_P \varepsilon}^{\pi+\delta} \cos(\alpha_P + \beta_0(s) - s)\, ds \\
&- \varepsilon \int_{\pi+\delta}^{2\pi-\alpha_P} \cos(-\alpha_P + \beta_1(s) - s)\, ds
\end{aligned}$$

where δ and $\beta_i(s)$ account for the circle not intersecting the skeleton midway and the distance function gradient field not being strictly a step function on $C_\varepsilon^P - arc_{\alpha_P}$. Therefore,

$$\begin{aligned}
\lim_{\varepsilon \to 0} \frac{\mathcal{F}_\varepsilon(P)}{2\pi\varepsilon} = &- \lim_{\varepsilon \to 0} \frac{\varepsilon}{2\pi\varepsilon} \int_{-\alpha_P}^{\alpha_P} ds \\
&- \lim_{\varepsilon \to 0} \frac{\varepsilon}{2\pi\varepsilon} \int_{\alpha_P}^{\pi+\delta} \cos(\alpha_P + \beta_0(s) - s)\, ds \\
&- \lim_{\varepsilon \to 0} \frac{\varepsilon}{2\pi\varepsilon} \int_{\pi+\delta}^{2\pi-\alpha_P} \cos(-\alpha_P + \beta_1(s) - s)\, ds \\
= &- \tfrac{1}{\pi} (\alpha_P + \sin \alpha_P)
\end{aligned}$$

since, as $\varepsilon \to 0$, δ, $\beta_0(s)$ and $\beta_1(s)$ vanish. Notice, however, that $\alpha_P = 0$ if the end-point is generated from a contour segment where the curvature is continuous.

3.4.3 Skeletal Junction and Pseudo-Junction Points

Let P be a skeletal junction point; that is where n skeletal curves meet. Let these curves be given by parameterizations $S_i(t)$ so that $S_i(0) = P$. Consider a circle of

[4]Here C_ε^P is as defined in Equation 3.5 but $\mathbf{t}_P = \lim_{\varepsilon \to 0} \mathbf{t}_{Q_\varepsilon}$.

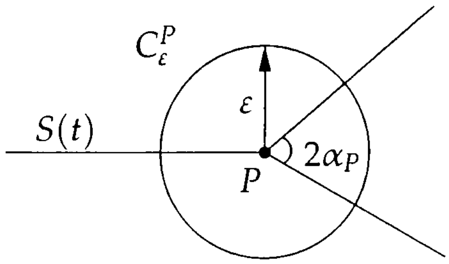

FIGURE 3.4. A circular neighborhood of radius ε around the end-point P. Along the arc of angle $2\alpha_P$ the gradient vectors agree (in orientation) with the normals to C_ε^P. Along the arc "above" $S(t)$ the gradient vectors all form an angle of α_P with $S'(0) = \mathbf{t}_P$. Similarly, for the arc "below," this angle is $-\alpha_P$.

radius ε centered at P. Denote it C_ε^P. For small enough ε, C_ε^P intersects the skeleton at precisely n regular points. Refer to them as $Q_\varepsilon^i = S_i(\varepsilon)$. Hence, to each there is a corresponding object angle. Define α_i as

$$\alpha_i = \lim_{\varepsilon \to 0} \alpha_{Q_\varepsilon^i}.$$

Now consider Figure 3.5 TOP. It suggests that $\sum_i 2\alpha_i = 2\pi$ for shapes without any circular arcs on their boundaries (such junction points will be referred to as *simple junction points*). Indeed, α_i is the angle between $S_i'(0)$ [5] and the line joining P to some point in P_C. To compute the outward flux through C_ε^P, we can divide the circle into n arcs, each corresponding to a skeletal curve. In particular, for $S_i(t)$ this would be the arc of angle $2\alpha_i$—denote these arcs as γ_i. For example, in Figure 3.5, the arc corresponding to $S_1(t)$ (i.e. γ_1) is the union of the two arcs of angle α_1. Notice that the distance function gradient field along γ_i behaves like that of a regular skeletal point with object angle α_i. Hence, the outward flux through it is

$$\begin{aligned}\mathcal{F}_{\gamma_i} = & -\varepsilon \int_{\delta_i}^{\alpha_i} \cos(\alpha_i + \beta_{0,i}(s) - s)\, ds \\ & -\varepsilon \int_{-\alpha_i}^{\delta_i} \cos(-\alpha_i + \beta_{1,i}(s) - s)\, ds\end{aligned}$$

where δ_i, $\beta_{0,i}(s)$ and $\beta_{1,i}(s)$ all vanish as $\varepsilon \to 0$. Thus, the total outward flux is

[5] Here $S_i'(0) = \lim_{t \to 0_+} S_i'(t)$.

Acknowledgments

I would first like to thank my supervisor, Prof. Kaleem Siddiqi, for his continued support. He has not only provided me with financial assistance, but has been a tremendous source of inspiration intellectually. He has nurtured without pushing emphasizing global understanding over detail.

I am grateful to my colleagues from Prof. Siddiqi's "Shape Group": Sylvain Bouix, Maxime Descoteaux, Abeer Ghuneim, Carlos Phillips and Alexander Vasilevskiy. You have all given me much to think about, especially the two Frenchmen: Sylvain and Carlos. Sylvain has always been willing to help either by providing a unique perspective or by simply contributing an interesting news flash. And Carlos, well, after many a street corner conversations and long discussions on politics, philosophy, society, science and the decidedly favorite subject of mathematics, my view of things has possibly changed ... but I'll never admit to that.

Finally, I gratefully acknowledge the financial support provided by the National Sciences and Engineering Research Council of Canada in the form of a PGS A scholarship.

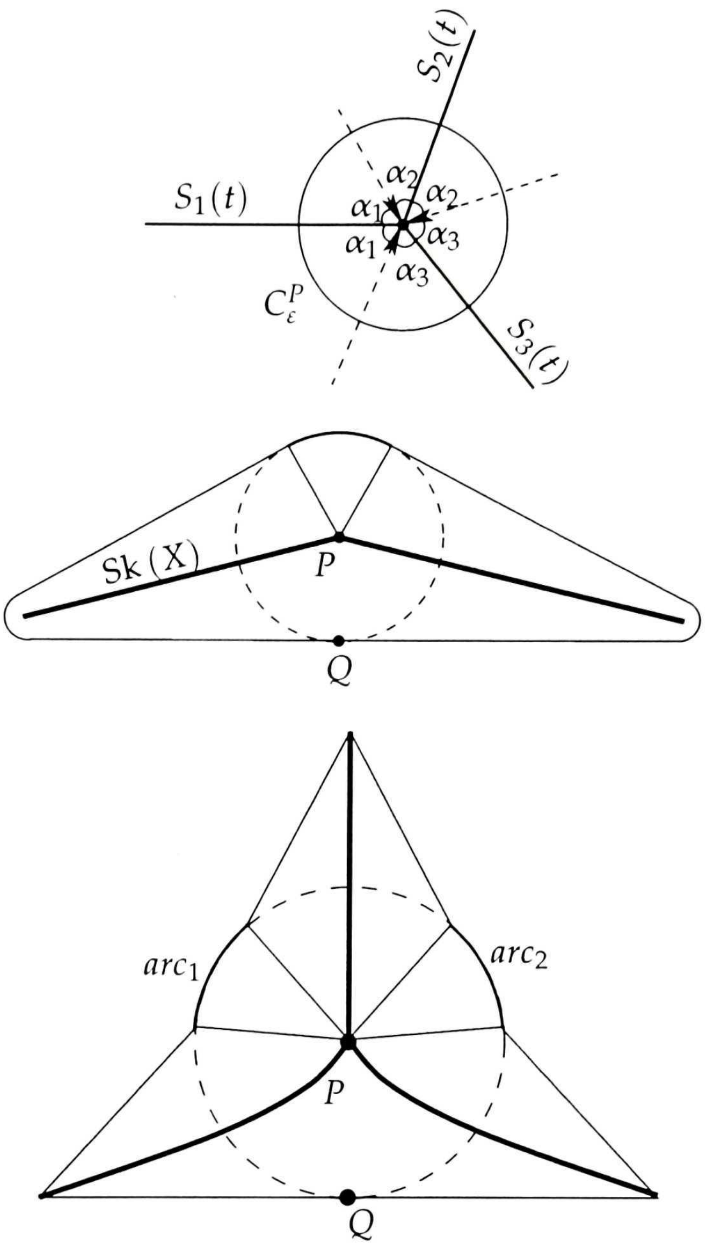

FIGURE 3.5. TOP: A circular neighborhood of radius ε around the simple junction point P. There are three skeletal curves denoted by $S_1(t)$, $S_2(t)$ and $S_3(t)$ respectively. The dashed lines link P and its closest points on the boundary (i.e. points in P_C). Note that $\alpha_1 + \alpha_2 + \alpha_3 = \pi$. CENTER: Pseudo-junction point. BOTTOM: General junction point.

$\mathcal{F}_\varepsilon(P) = \sum_{i=1}^{n} \mathcal{F}_{\gamma_i}$ and the average outward flux becomes

$$\begin{aligned}
\lim_{\varepsilon \to 0} \frac{\mathcal{F}_\varepsilon(P)}{2\pi\varepsilon} &= \sum_{i=1}^{n} \lim_{\varepsilon \to 0} \mathcal{F}_{\gamma_i} \\
&= -\frac{1}{\pi} \sum_{i=1}^{n} \sin \alpha_i .
\end{aligned}$$

Now, for general junction points, we must also consider the effects circular arcs on the boundary of the shape (see Figure 3.5 BOTTOM). Naturally, in such situations, the sum of object angles from the n skeletal curves does not add to 2π, we must also include the angles θ_j of the arcs arc_j from the boundary. Hence,

$$\mathcal{F}_\varepsilon(P) = \sum_{i=1}^{n} \mathcal{F}_{\gamma_i} + \sum_{j=1}^{k} \mathcal{F}_{arc_j}$$

where $\mathcal{F}_{arc_j}$ is $-\varepsilon\theta_i$, n is the number of skeletal curves and k is the number of circular arcs on the boundary. Notice that pseudo-junction points are those where $n = 2$ and $k = 1$, thus the general result for all types of junction points is

$$\mathscr{F}_\varepsilon(P) = -\frac{1}{\pi} \sum_{i=1}^{n} \sin \alpha_i - \frac{1}{2\pi} \sum_{j=1}^{k} \theta_j \,. \tag{3.8}$$

3.4.4 Non-Skeletal Points

Now, let P be a non-skeletal point. In particular, there exists an ε small enough, so that C_ε^P contains no skeletal points. Hence, the distance function gradient field along the circle is continuous. Thus, we can write

$$\mathcal{F}_\varepsilon(P) = \varepsilon \int_0^{2\pi} \cos(\alpha + \beta(s) - s)\, ds,$$

where α is any orientation of the distance function gradient field along C_ε^P and $\lim_{\varepsilon\to 0} \sup_{s\in[0,2\pi]} |\beta(s)| = 0$. Hence,

$$\lim_{\varepsilon\to 0} \frac{\mathcal{F}_\varepsilon(P)}{2\pi\varepsilon} = \frac{1}{2\pi} \int_0^{2\pi} \cos(\alpha - s)\, ds = 0.$$

3.5 Summary

Following the formal introduction of shape and skeleton in Chapter 2, we developed in this chapter a criterion that classifies points in the interior of a shape as either belonging to the skeleton or not. We extended the Divergence Theorem and used this new result to show that the limit of the average outward flux through the boundary of a region as it shrinks to a point may be used to distinguish between

POINT TYPE	$\lim_{\varepsilon \to 0} \frac{\mathcal{F}_\varepsilon(P)}{2\pi\varepsilon}$
Regular Points (Figure 3.2)	$-\frac{2}{\pi}\sin\alpha$
End-Points (Figure 3.4)	$-\frac{1}{\pi}(\sin\alpha + \alpha)$
Junction Points (Figure 3.5)	$-\frac{1}{\pi}\sum_{i=1}^{n}\sin\alpha_i - \frac{1}{2\pi}\sum_{j=1}^{k}\theta_j$
Non-Skeletal Points	0

TABLE 3.1. A summary of results relating the limit values of the average outward flux to the object angle for shrinking circular neighborhoods. Note that for contours of type C^3 (i.e. with continuous curvature), α will be zero for the case of end-points.

regular skeletal points and non-skeletal points. Our analysis also provides a necessary and sufficient condition that must be satisfied by any limiting process so that AOF criterion may be applied. Further, for the special case of circular regions, we showed that the criterion also covered the extreme points of skeletal curves; that is, as the circle shrinks to a point P, the AOF tends to zero if and only if $P \notin \mathrm{Sk}(X)$. We also saw that the nonzero limit value of the AOF on regular $P \in \mathrm{Sk}(X)$, is actually the object angle in disguise. Table 3.1 summarizes the exact behavior.

These results suggest a computational approach for the extraction of $\mathrm{Sk}(X)$ given a shape X: keep only the points P where the approximation to $\lim_{R \to P} \mathscr{F}_{\partial R}$ is strictly negative. However, there are several issues with this direct algorithm that must be addressed in order to obtain a robust and accurate approximation of $\mathrm{Sk}(X)$; this is the topic of the next chapter.

Chapter 4

Flux-Based Skeletonization

In Chapter 3, we derived an average outward flux criterion that allows us to distinguish between points on the skeleton of a shape and those which do not belong to that structure. However, the discussion was carried out in the continuous domain $\mathbb{R}^2$ taking advantage of a tool not present in discrete computations: limits. A digital machine may only approximate limits so, if we are to use the theory from Chapter 3, we must understand the limiting behavior of the AOF even better. In this chapter, we shall analyze the distance function gradient field on a discrete lattice in order to adapt the continuous criterion to a square grid. This will then allow us to develop a procedure that computes a good discrete approximation of a shape's skeleton: the approximating points will be arbitrarily close to the real structure, ordered along skeletal curves, and we shall be able to guarantee a user specified maximal separation between consecutive points.

In Section 4.1, we shall introduce the computational setup and discuss the meaning of homotopy equivalence on a discrete grid. Then, in Section 4.2, we shall derive an algorithm that produces a thin discrete approximation to the skeleton no more than a grid spacing away from the real skeleton. We shall also see there how the output of this procedure can be converted into a graph as discussed in Chapter 2. In Section 4.3, we shall analyze the adequacy of this approximation by studying the effects of thresholding the AOF criterion. Then, in Section 4.4, we shall see how the points in this approximation can be shifted to be within a user-specified distance to the real structure and, just before the summary (Section 4.6), we shall perform, in Section 4.5, a number of numerical experiments and corroborate our theory.

4.1 Computational Setup

Given a shape X, its contour curves may be represented discretely through splines. For the purposes of computer vision, this does not seem to be a restriction and so it will be the representation of choice in what follows. In fact, as far as the sample implementation is concerned, contour curves will be assumed to be a collection of line segments satisfying the conditions of Definition 2.6. In Section 4.1.1, we give the details of acceptable representations for the computational framework in this chapter and, in Section 4.1.2, we discuss the notion of homotopy thinning on a square grid, as well as reasonable parameters for the estimation of limits there.

4.1.1 ∂X, ∇D_X and Discrete AOF

In order to use the AOF criterion developed earlier, the computational setup should make it easy to determine the value of the distance function at any point on $\mathbb{R}^2$, as well as its gradient. If those quantities were exact (as accurate as hardware permits), then the approximation of the average outward flux through any circle of nontrivial radius could be computed within any specified tolerance of the real value. With these goals in mind, we now turn to the question of how to represent shapes efficiently.

The curves forming ∂X (see Definition 2.6) can be given a notion of inside consistent with X. Since any such curve is simple and closed, by definition exactly one of the two regions it separates contains the shape — that is the *inside* of the curve. Now, if a point $P \in \mathbb{R}^2$ is closest to a contour curve $C(t)$, then it is inside the shape if and only if it is in the inside region for $C(t)$. Hence, one can define a signed distance function D_C for each $C(t)$:

$$D_{C_i}(P) = \chi_{C_i}(P) \inf_{Q \in C_i(t)} \mathrm{d}\,(P,\, Q)$$

where $\chi_{C_i}(P)$ is -1 whenever P is in the inside region of $C_i(t)$ and 1 otherwise. Thus, the collection of D_{C_i} "split" D_X, that is

$$D_X(P) = \min_{i=1}^{n} D_{C_i}(P).$$

Notice that, by Theorem 2.9(a), ∇D_X also splits.

Therefore, the object oriented paradigm may be used to perform the calcula-

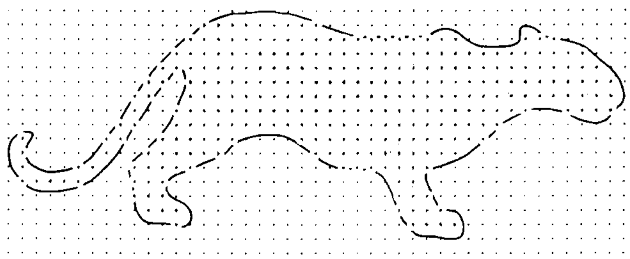

FIGURE 4.1. The lattice points overlayed on top of a shape defined in $\mathbb{R}^2$. These are the locations where the average outward flux is computed and where the thinning takes place. The gray circles identify the points in $L_\sigma(X)$.

tions. Think of a shape as an object which collects several curves. If, in turn, each curve is an object endowed with a signed distance function as above and ∇D_C computed as suggested by the theorem, then the shape object knows how to compute D_X and ∇D_X. To obtain D_C, a curve object may be seen as a set of atomic curve segments each endowed with the above distance function collected exactly similarly. Thus, if each atomic segment can provide exact information (as is the case for line segments used in the sample implementation), then exact information is returned by D_X and ∇D_X.

Unfortunately, the average outward flux must be numerically approximated. Recall that the average outward flux through a circle of radius r centered at P is given by

$$\mathscr{F}_r(P) = \frac{\int \langle\, \nabla D_X(P + r\mathcal{N}(s)),\ \mathcal{N}(s)\,\rangle\, \mathrm{d}s}{2\pi r}$$

where $\mathcal{N}(s)$ is the outward normal to the circle and the integral is over the circle. Applying the Trapezoidal rule, the numerator is approximated by

$$\int \langle\, \nabla D_X(P + r\mathcal{N}(s)),\ \mathcal{N}(s)\,\rangle\, \mathrm{d}s = \frac{2\pi r}{n} \sum_{i=0}^{n-1} \langle\, \nabla D(P + r\mathcal{N}(i)),\ \mathcal{N}(i)\,\rangle \qquad (4.1)$$

where $\mathcal{N}(i) = (\cos(i2\pi/n), \sin(i2\pi/n))$ is the outward normal to the circle at the sampled locations. To obtain the average outward flux, one simply divides the above by $2\pi r$.

4.1.2 The Discrete Lattice $\mathbb{L}_\sigma$

The skeleton of a shape will be approximated by a finite set of points chosen from a discrete grid—the square lattice $\mathbb{L}_\sigma$. The representations in the previous subsection did not assume anything about where on the $\mathbb{R}^2$ plane the calculations are to be carried out. Consequently, there is no restriction on the size and position of the grid and therefore, as will become evident in what follows, the quality of the skeleton approximation is only limited by the choice of σ. But first, the necessary definitions are discussed as well as how the discrete computations translate into properties of shapes in $\mathbb{R}^2$.

Let us begin by formally introducing the discrete lattice. It is denoted $\mathbb{L}_\sigma$ and, as a set, can be thought of as the "square grid" in $\mathbb{R}^2$ with *spacing*[1] σ, i.e.

$$\mathbb{L}_\sigma = \sigma\mathbb{Z} \times \sigma\mathbb{Z} \subset \mathbb{R}^2.$$

Now, if X is a shape, let $L_\sigma(X)$ be the discrete representation of X, i.e.

$$L_\sigma(X) = \{P \in \mathbb{L}_\sigma \; : \; P \in X\}.$$

A discrete approximation to the skeleton of X, denoted $\mathrm{Sk}_\sigma(X)$, will be computed from $L_\sigma(X)$. The idea is to thin $L_\sigma(X)$ preserving homotopy type and using the discrete average outward flux as a stopping criterion.

However, homotopy thinning in the discrete setting must be related to the continuous case; in particular, it is desirable to somehow assign homotopy type to $\mathrm{Sk}_\sigma(X)$ and be able to verify that it is the same as that of X.[2] To do this, let us introduce an operator, $\mathcal{S}_\sigma(U)$ (i.e. "the shape of U"), that assigns a "shape" to a discrete set of points $U \subseteq \mathbb{L}_\sigma$. Define it as

$$\mathcal{S}_\sigma(U) = \bigcup_{P \in U} \overline{\mathcal{B}_{2\sigma/3}(P)};$$

that is, as the union of all closed balls of radius $2\sigma/3$ centered at a point in U. The radius is such that it is possible to find a "small enough" σ for which $\mathcal{S}_\sigma(L_\sigma(X))$ is of the same homotopy type as X, i.e. the following holds (see Appendix B, Proposition B.1 for the proof)

[1] or, equivalently, *resolution*.

[2] Meyer [34] provides an, essentially, equivalent solution to the one presented here, but he assumes shapes to be discrete entities.

Proposition 4.1. *Given a shape X with no isolated points in ∂X, then there must be $\sigma_0 > 0$ such that if $\sigma < \sigma_0$ then $\mathcal{S}_\sigma(L_\sigma(X))$ is homotopy equivalent to X.*

Hence, a *homotopy thinning* is performed on $L_\sigma(X)$ if and only if, at any step of the thinning, the resulting set U is such that $\mathcal{S}_\sigma(U)$ is homotopy equivalent to $\mathcal{S}_\sigma(L_\sigma(X))$. If σ is small enough so that $\mathcal{S}_\sigma(L_\sigma(X))$ is homotopy equivalent to X and $\mathrm{Sk}_\sigma(X)$ is the result of a homotopy thinning on $L_\sigma(X)$, then $\mathcal{S}_\sigma(\mathrm{Sk}_\sigma(X))$ is homotopy equivalent to X.

Another important consideration in computing $\mathrm{Sk}_\sigma(X)$ is the calculation of the discrete average outward flux. The calculation in Equation 4.1 involves two parameters: (1) the number of sample points and (2) the radius of the circle. As will be shown later, (1) affects the approximation to the actual average flux and error bounds are available, so let us look at (2).

Chapter 3 showed the behavior of the average outward flux for very small values of r (indeed, as $r \to 0$). Unfortunately, on $\mathbb{L}_\sigma$, r may not be zero and if it is too small (e.g. $r = 0.1\sigma$) then evaluating the AOF at some lattice points near the skeleton might not be affected by the skeleton at all and, consequently, would be an inadequate criterion for detecting such points. On the other hand, choosing a large r (e.g. 2σ) would make the AOF a poor skeleton detector as it would "blur" the information and hinder localization. A good compromise, then, (one for which properties are shown later) is

$$r = \sigma.$$

This choice does not sacrifice localization for detection and it can also handle an n-point approximation of the AOF (see Equation 4.1). The following section elaborates.

4.2 The Algorithm

Chapter 3 provided a continuous criterion—the limit of the AOF through a shrinking circle—for separating skeletal points from non-skeletal points in $\mathbb{R}^2$ and, as we have just seen, it is possible to approximate it on a discrete lattice by evaluating the average outward flux through a small circle (radius $r = \sigma$). However, it is not immediately clear how this approximation may be used to, in turn, approximate the skeleton on $L_\sigma(X)$ and obtain a thin structure, which is topologically equivalent to the original shape. We shall see, in Section 4.2.1, that the most straightforward approach—simply thresholding on the AOF—is not an acceptable solution: the result is thick and possibly disconnected. The second problem may be fixed by a homotopy thinning, a class of algorithms discussed in Section 4.2.2, but we'll see, in Section 4.2.3, that the first problem needs more work. In Section 4.2.4, we'll discuss the selection of discrete end-points, which will then be used to thin again and produce the desired result. Finally, in Section 4.2.5, we'll see how to obtain the graph interpretation (discussed in Chapter 2) of the thin skeleton.

4.2.1 Thresholding the AOF on $\mathbb{L}_\sigma$

Chapter 3 presented a criterion that distinguishes between points on the skeleton of a shape from those which are not on the skeleton—in the limit as a circle shrinks to a point, the average outward flux is nonzero in the former case and identically zero in the latter. Chapter 2 on the other hand, discussed properties of the signed distance function D_X induced by a 2D shape X, and showed that ∇D_X is continuous almost everywhere. It is therefore reasonable to expect that the behavior of the AOF through circles of small radius be similar to the limiting case; in other words, it should be possible to obtain a discrete approximation of the skeleton, $\mathrm{Sk}_\sigma(X)$, by simply keeping points in $L_\sigma(X)$ for which the AOF is above a certain value. Of course, the quality of this approximation depends on how the AOF is obtained and on the threshold.

Although it is difficult to make general statements about the behavior of the AOF through a circle of radius σ, it is possible to bound the AOF if the circle C_σ^P is completely inside the shape and C_σ^P does not contain any skeletal points in its interior. The bound is a function of maximum variation of the orientation of ∇D_X; i.e. (see Lemma B.2 on page 84 for the proof)

Lemma 4.2 (Deviation). *Let $E(s) : [\theta_0, \theta_1] \to [-\delta/2, \delta/2]$ for $0 \le \delta \le 2\pi$. Then*

$$\left| \int_0^{2\pi} \cos(\alpha + E(s) + s)\ \mathrm{d}s \right| \le \delta + 2\sin\left(\frac{\delta}{2}\right)$$

Further, if the integral is numerically approximated by the Trapezoidal rule with n sample points and separation $\frac{2\pi}{n}$ between them, then

$$\left| I^{(n)} \right| \le \delta + 2\sin\left(\frac{\delta}{2}\right) + \mathrm{NumErr}(n)\ .$$

where the top integral actually denotes the outward flux through a circle (recall Section 3.4) and $I^{(n)}$ denotes the n-point approximation given by Equation 4.1 with numerical error $\mathrm{NumErr}(n)$. So, if δ can be reasonably estimated for circles as above, the lemma provides a bound for the continuous version of the AOF through those circles (or if $\mathrm{NumErr}(n) = 0$). The following lemma provides a tight bound on δ (refer to Lemma B.3 on page 87 for a proof):

Lemma 4.3. *Let X be a shape and ∂X its boundary. Let C_σ^P be a circle centered at P with radius σ. Assume that C_σ^P does not contain in its interior any skeletal points or points on ∂X. Let δ be the maximum difference in orientation of ∇D_X through C_σ^P. Then,*

$$\delta \leq 2 \arcsin\left(\frac{\sigma}{d+\sigma}\right)$$

where $d = \min\left\{\mathrm{d}\left(C_\sigma^P, \partial X\right), \mathrm{d}\left(C_\sigma^P, \mathrm{Sk}\left(X\right)\right)\right\} > 0.$

Hence, collecting Lemma 4.2 and Lemma 4.3, the following can be said about the numerical approximations of the average outward flux under the assumptions of Lemma 4.3:

Corollary 4.4. *Let X, ∂X, C_σ^P and d be as in Lemma 4.3 and suppose the average outward flux through C_σ^P is approximated by n sample points as in Lemma 4.2. If $\mathscr{F}_\sigma^{(n)}(P)$ denotes this approximation, then*

$$\left|\mathscr{F}_\sigma^{(n)}(P)\right| \leq \frac{1}{2\pi}\left[2 \arcsin\left(\frac{\sigma}{d+\sigma}\right) + 2\left(\frac{\sigma}{d+\sigma}\right) + \mathrm{NumErr}(n)\right].$$

Therefore, if the numerical error could be made negligible, then thresholding according to the above bound would guarantee that the points kept are at most distance d away from the skeleton. The following result suggests how to achieve this (see Lemma B.4 on page 88).

Lemma 4.5. *Let X be a shape and denote by $\kappa(t)$ the curvature function of ∂X wherever it is defined. If the boundary is made of line segments then*

$$\mathrm{NumErr}(n) \leq \frac{4}{3}\frac{\pi^3}{n^2}.$$

Suppose $X \subseteq \overline{\mathcal{B}_R(P)}$, $|\kappa(t)| \leq K_1 \in \mathbb{R}$ and $|\kappa'(t)| \leq K_2 \in \mathbb{R}$. If $\frac{1}{K_1} - \frac{R}{2} > \sigma$ and

Table of Contents

Abstract i

Résumé ii

Contribution of Authors iii

Acknowledgments iv

1 Introduction 1

1.1 Shape and Shape Representation 2

1.2 Computing Skeletons 5

1.2.1 Topological Thinning 6

1.2.2 Distance Map Methods 7

1.2.3 Grassfire Simulation and Wave Propagation 8

1.2.4 Voronoi Diagram Approaches 9

1.2.5 Computational Difficulties 10

1.3 Contributions 10

2 Definitions and Properties 12

2.1 Shapes 12

2.2 Euclidean Distance Function 15

2.3 Skeletons and Skeleton Transform 16

2.4 Boundary Representation through $\mathbf{ST}(X)$ 23

2.5 Summary 28

3 The AOF Criterion 29

3.1 Preliminaries 29

3.2 An Extension of the Divergence Theorem 30

3.3 Skeletal Point Selection Criterion 32

INPUT: $L_\sigma(X) \subset \mathbb{L}_\sigma$
OUTPUT: $S_{\mathrm{Thr}(d)}(X) \subset \mathbb{L}_\sigma$
$S_{\mathrm{Thr}(d)}(X) = \varnothing$
for all $P \in L_\sigma(X)$ **do**
  if $\mathscr{F}_\sigma^{(n)}(P) < -\mathrm{Thr}(d)$ **then**
    Insert P in $S_{\mathrm{Thr}(d)}(X)$
  end if
end for

ALGORITHM 1. Skeleton approximation by simple thresholding.

$n = 2^k$ *for* $k \geq 2$ *then*

$$\mathrm{NumErr}(n) \leq \frac{4}{3}\frac{\pi^3}{n^2}\mathcal{E}\left(\frac{1}{K_1} - \frac{R}{2}\right).$$

where

$$\begin{aligned}\mathcal{E}(\xi) &\leq K_2\frac{R}{\xi}\frac{8+38\xi^2+25\xi^3+8\xi^4+28\xi+\xi^5}{\xi^5} \\ &+ \frac{8+16\xi+2\xi^3+10\xi^2}{\xi^5}.\end{aligned}$$

Thus, we define a thresholding function as follows:

Definition 4.6. The *threshold function* $\mathrm{Thr}(d)$ is

$$\mathrm{Thr}(d) = \frac{1}{2\pi}\left[2\arcsin\left(\frac{\sigma}{d+\sigma}\right) + 2\left(\frac{\sigma}{d+\sigma}\right)\right].$$

Now the simple algorithm for approximating the skeleton of a shape consists of computing the AOF approximation at all points in $L_\sigma(X)$ and keeping only those that satisfy a threshold given by $\mathrm{Thr}(d)$; see Algorithm 1.

So far, we only have negative results about the performance of this procedure—Corollary 4.4 estimates which lattice points are *discarded* by thresholding according to $\mathrm{Thr}(d)$. Let us now examine which elements of $L_\sigma(X)$ that are close to $\mathrm{Sk}(X)$ might be kept. First, assume that the circle of radius σ centered at $P \in L_\sigma(X)$ is intersected by a curve segment of the skeleton. Without loss of generality, P is at most $\sigma/2$ away from the skeleton. In this case it is possible to amend Lemma 4.2 to obtain (Lemma B.2 on page 84 for the proof)

Lemma 4.7. *Let P be such that C_σ^P intersects the skeleton, i.e. the interior of the circle contains points in* $\mathrm{Sk}(X)$. *Then, the n-point approximation of the AOF is bounded by*

$$-\frac{1}{2\pi}(4\sin(\alpha)+E) \leq \mathscr{F}_\sigma^{(n)}(P) \leq -\frac{1}{2\pi}\left(2\sqrt{3}\sin(\alpha)-E\right)$$

where $E = 4\delta + \mathrm{NumErr}(n)$, *and* ∇D_X *along the circle meets the skeleton at an angle bounded by* $[\alpha-\delta/2, \alpha+\delta/2]$.

Now, using Algorithm 1 and the bound from Lemma 4.7, it is possible to estimate which points will be kept. Consider

$$-\frac{1}{2\pi}(4\sin(\alpha)+E) \leq -\frac{1}{2\pi}\left[2\arcsin\left(\frac{\sigma}{d+\sigma}\right)+2\left(\frac{\sigma}{d+\sigma}\right)\right]$$

so

$$\alpha \geq \arcsin\left[\frac{1}{4}\left(2\arcsin\left(\frac{\sigma}{d+\sigma}\right)+2\left(\frac{\sigma}{d+\sigma}\right)-E\right)\right] \tag{4.2}$$

and, exactly similarly,

$$\alpha \leq \arcsin\left[\frac{1}{2\sqrt{3}}\left(2\arcsin\left(\frac{\sigma}{d+\sigma}\right)+2\left(\frac{\sigma}{d+\sigma}\right)+E\right)\right]. \tag{4.3}$$

Therefore, recalling Corollary 4.4,

Theorem 4.8 (Threshold Lemma). *The points P in $L_\sigma(X)$ which satisfy*

$$\mathscr{F}_\sigma^{(n)}(P) \leq -\mathrm{Thr}(d)$$

are

1. *No further than d from the skeleton (for large enough n); and*
2. *Those which are closer than $\sigma/2$ from the skeleton are such that the object angle of the skeletal points contained in a σ-ball around the point is bounded by*

$$\arcsin\left[\frac{1}{4}(2\pi\mathrm{Thr}(d)-E)\right] \leq \alpha \leq \arcsin\left[\frac{1}{2\sqrt{3}}(2\pi\mathrm{Thr}(d)+E)\right] \tag{4.4}$$

where $E = 4\delta + \mathrm{NumErr}(n)$, *and* ∇D_X *along the circle* C_σ^P *meets the skeleton at an angle bounded by* $[\alpha-\delta/2, \alpha+\delta/2]$.

Notice that the lattice points that would best approximate the skeleton are those in part 2 of the theorem; that is, points no further than $\sigma/2$ from the $\mathrm{Sk}(X)$. Indeed, it is desirable to obtain

$$\mathrm{Sk}_{\sigma}(X) = \{P \in \mathbb{L}_{\sigma} : \mathrm{d}(P, \mathrm{Sk}(X)) \leq \sigma/2\}. \tag{4.5}$$

However, simple thresholding does not provide this ideal approximation; in fact, serious issues arise due to the global nature of the operation. To illustrate, let us assume that the error term E in Theorem 4.8 is negligible. Then, the points kept by a threshold $\mathrm{Thr}(d)$ will "be of object angle"[3] which is a function of $\mathrm{Thr}(d)$. At the same time, there might be $P \in \mathbb{L}_{\sigma}$ up to d away from the real skeleton that will also be kept by this procedure. Now, if $S_{\mathrm{Thr}(d)}(X)$ denotes the approximation of the skeleton by thresholding, a real problem arises if $\mathcal{S}_{\sigma}\left(S_{\mathrm{Thr}(d)}(X)\right)$ is not of the same homotopy type as $\mathcal{S}_{\sigma}(L_{\sigma}(X))$, which, regrettably, is quite possible.

Consider the examples in Figure 4.2; the shape is shown on top and the contents of $S_{\mathrm{Thr}(d)}(X)$ for $d = 3$ just below. The flux was approximated with $n = 30$ and it is easily seen that δ is zero at all points of the skeleton except on the line in the middle where it is proportional to $\frac{\sin\alpha}{r}$ (r is distance to the contour). The object angle near the junction points on the side closer to the center is lowest relative to the rest of the skeleton which explains the lack of approximating pixels there. On the other hand, the threshold cannot remove the "noise pixels" near the end-points corresponding to the centers of curvature of the two circular arcs. Note that this situation is inescapable: if a threshold is picked high enough to remove some of the noise at the end-points (e.g. $d = 1$), then the object angle of the points kept by this threshold will increase and the gaps in $S_{\mathrm{Thr}(d)}(X)$ will become even wider.

This example demonstrates that, even for generic shapes, any α-skeleton may have regular points with object angles much lower than α. Thus, the result of thresholding on the object angle will possibly be disconnected or trivial (empty). Hence, Algorithm 1 will not produce a thin—as in Equation 4.5—skeleton approximation which is topologically equivalent to the initial shape. In the following, we shall couple thresholding with a homotopy thinning approach.

4.2.2 Homotopy Thinning

Section 4.1.2 introduced the notion of a homotopy thinning on subsets of $\mathbb{L}_{\sigma}$ but did not discuss any such procedure. Vincent [54] discusses such an algorithm and provides removability tables. However, it turns out that there is a simple nec-

[3] In the sense of Equation 4.4.

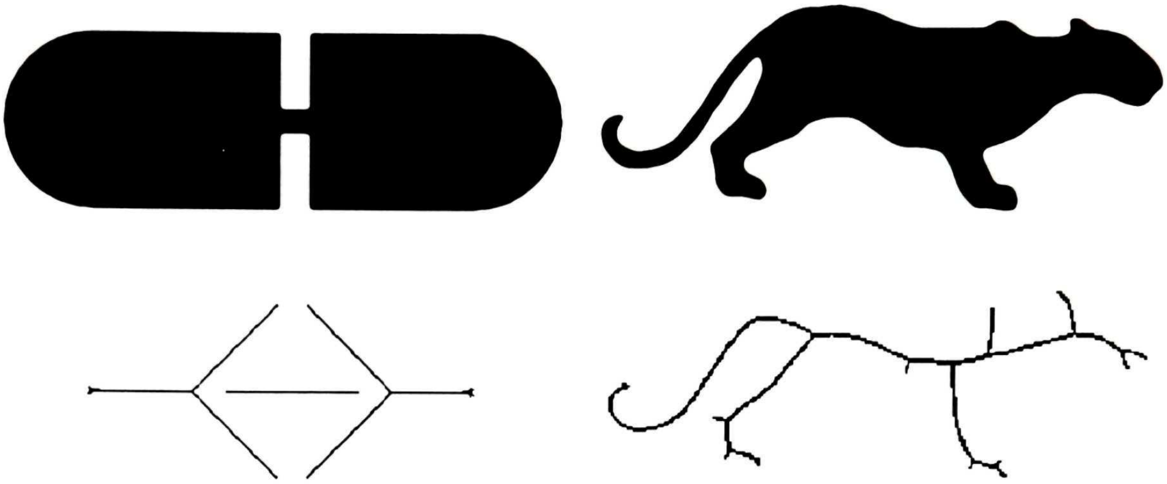

$d = 3\sigma$

$d = \sigma$

FIGURE 4.2. Simple thresholding. Thresholding for a thin skeleton may disconnect the approximation.

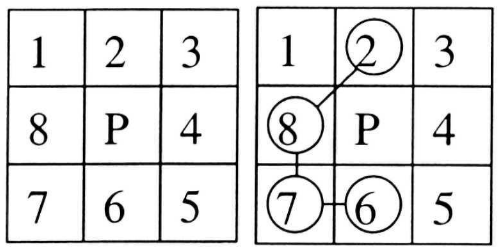

FIGURE 4.3. LEFT: A 3x3 neighborhood of a 2D digital point P in a rectangular lattice. RIGHT: An example neighborhood graph for which P can be removed. Note that there is no edge between neighbors 6 and 8 (see text).

1: INPUT: $U \subset \mathbb{L}_\sigma$
2: OUTPUT: $U' \subset U$
3: **while** There is a removable $P \in \partial U$ **do**
4: Choose $P \in \partial U$
5: **if** P is removable **then**
6: $U = U - \{P\}$
7: **end if**
8: **end while**

ALGORITHM 2. Generic thinning. Here ∂U is the discrete boundary of U which, of course, assumes that U is bounded. Step 4 would incorporate the criterion in Proposition 4.9 to turn this procedure into a homotopy thinning.

essary and sufficient condition that determines if a thinning scheme is homotopy preserving or not (see Dimitrov *et al.* [18]).

We begin with the notion of thinning. Algorithm 2 is the description of a generic thinning procedure: at each step somehow select a point on the boundary of the remaining object and remove it if it makes sense for the application at hand. Typically, an order is assigned to the pixels on ∂U and removal is attempted accordingly; the algorithm terminates if all candidates must be kept. However, removability plays an even more important role—the topology (on $\mathcal{S}_\sigma(U)$, say) induced by U can be directly controlled by tuning this criterion.

Let $U \subset \mathbb{L}_\sigma$ be the initial (bounded) set and denote by $U^{(n)}$ the resulting set after removing n points from U. Let P be such that $U^{(n-1)} = U^{(n)} \cup \{P\}$. The homotopy type of $\mathcal{S}_\sigma\left(U^{(n-1)}\right)$ and $\mathcal{S}_\sigma\left(U^{(n)}\right)$ will be different if and only if the removal of P creates a hole in $\mathcal{S}_\sigma\left(U^{(n-1)}\right)$ or locally (as in Figure 4.3 left) disconnects it. It is convenient to view this as a graph problem. Consider the 3x3 neighborhood of P as shown in Figure 4.3, and select those neighbors that are also in $U^{(n-1)}$. Construct a neighborhood graph by placing edges between pairs of points that are 4-adjacent or 8-adjacent to one another. If any of the 3-tuples $\{2,3,4\}$, $\{4,5,6\}$, $\{6,7,8\}$ or $\{8,1,2\}$ are nodes of the graph, remove the corresponding diagonal edges $\{2,4\}$, $\{4,6\}$, $\{6,8\}$ or $\{8,2\}$, respectively. This ensures that there are no degenerate cycles in the neighborhood graph (cycles of length 3). Now, observe that if the removal of P disconnects $U^{(n-1)}$ or introduces a hole, the neighborhood graph will not be connected or will have a cycle, respectively. Conversely, a connected graph that has no cycles (i.e. a tree) means that $U^{(n)}$ will be of the same homotopy type as $U^{(n-1)}$. Hence, the criterion may be expressed as follows:

Proposition 4.9. *$\mathcal{S}_\sigma\left(U^{(n-1)}\right) = \mathcal{S}_\sigma\left(U^{(n)} \cup \{P\}\right)$ and $\mathcal{S}_\sigma\left(U^{(n)}\right)$ are homotopy equivalent if and only if the neighborhood graph of P, with cycles of length 3 removed, is a tree.*

The application of this criterion is fairly simple: a graph is a tree if and only if its Euler characteristic number (i.e. number of nodes minus number of edges) is identical to 1.

Hence, the first problem identified in the previous subsection—that of preserving topology—may be solved by combining thresholding and Proposition 4.9; which is the approach discussed next.

4.2.3 Thick Skeletons

In Section 4.2.1 a first attempt was made to approximate the skeleton of a shape but it was demonstrated to be flawed: simply keeping lattice points with high enough flux, according to a threshold, could (and, by and large, always does) disconnect the skeleton thereby yielding an approximation which is not topologically equivalent to the original shape. The previous subsection on the other hand, suggests an approach to fix this problem but, as will become clear, the final issue pertaining to the thickness of the approximation will have to be addressed separately.

Consider the condition on line 4 in Algorithm 2. Combining thresholding with homotopy thinning amounts to defining this criterion and specifying how to choose a point on the boundary in line 3. Therefore, let a point $P \in \partial U$ be removable if and only if $\mathscr{F}_\sigma^{(n)}(P) > -\mathrm{Thr}(d)$ and its removal maintains the topology unchanged, i.e. using Proposition 4.9. Now, in order to obtain an approximation which is as close as possible to $\mathrm{Sk}(X)$, the points on the boundary that should be removed first are those furthest from $\mathrm{Sk}(X)$ or, equivalently, those closest to ∂X. Algorithm 3 is the ensuing procedure and Figure 4.4 provides some examples of its application.

There is an important relationship between the output of Algorithm 3 and that of Algorithm 1. Let $S_{Thick(d)}(X)$ denote the approximation of $\mathrm{Sk}(X)$ obtained by Algorithm 3 and let $S_{\mathrm{Thr}(d)}(X)$, as before, be the approximation obtained by simply thresholding. Then

$$S_{\mathrm{Thr}(d)}(X) \subseteq S_{Thick(d)}(X)\,;$$

in other words, this procedure may only add points to $S_{\mathrm{Thr}(d)}(X)$ and if it does, then those will be as close as possible to $\mathrm{Sk}(d)$, i.e. less than $\sigma/2$ from it. This

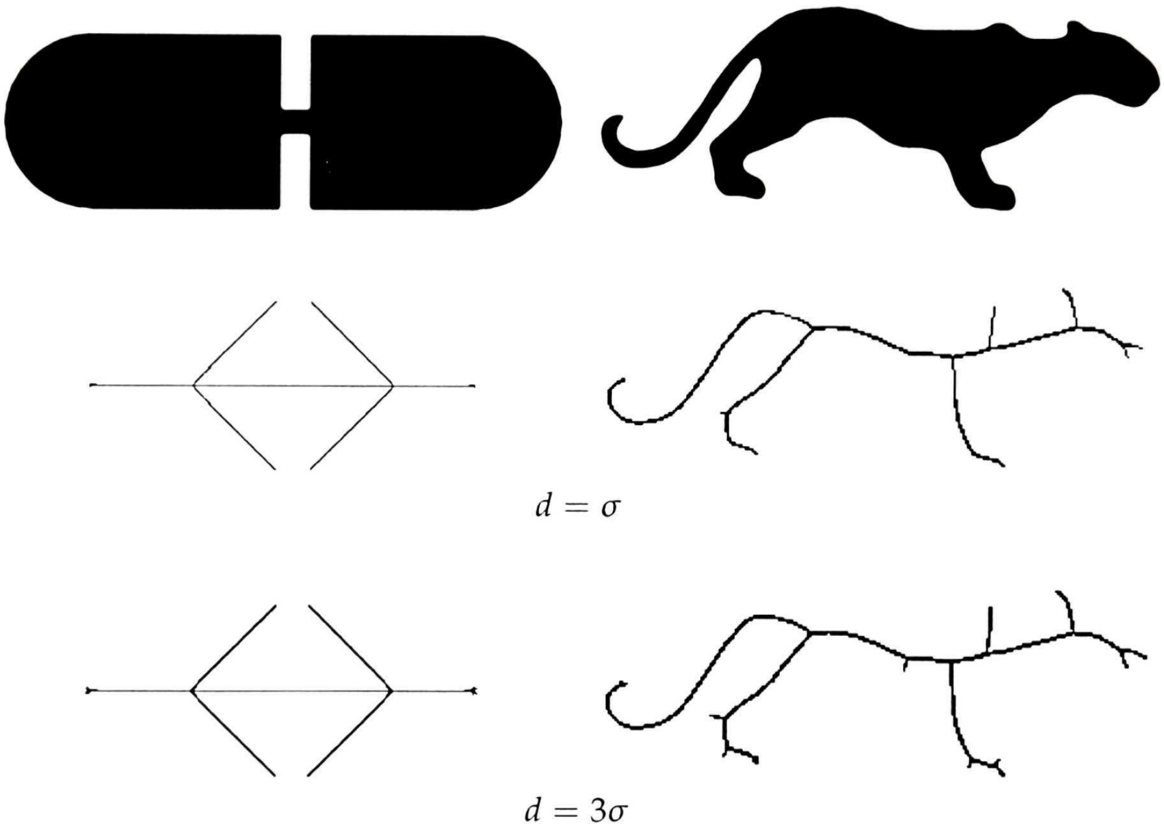

FIGURE 4.4. Thick skeletons. Contrast with Figure 4.2; homotopy thinning maintains connectivity.

INPUT: $U = L_\sigma(X) \subset \mathbb{L}_\sigma$
OUTPUT: $S_{Thick(d)}(X) \subset L_\sigma(X)$
while There is a removable $P \in \partial U$ **do**
  Choose $P \in \partial U$ closest to ∂X among remaining choices.
  if $\mathscr{F}_\sigma^{(n)}(P) > -\mathrm{Thr}(d)$ **and** $NG(P)$ is a tree **then**
    $U = U - \{P\}$
  end if
end while
$S_{Thick(d)}(X) = U$

ALGORITHM 3. Thick skeleton algorithm. $NG(P)$ denotes the neighborhood graph of P as described in Section 4.2.2.

is a consequence of the ordering of the points on ∂U in Algorithm 3; a point $P \in S_{Thick(d)}(X)$ which is not in $S_{\mathrm{Thr}(d)}(X)$ must be furthest possible from the boundary among its neighbors that have not been kept.

However, the relationship also implies that $S_{Thick(d)}(X)$ can possibly contain lattice points as far as d from the real skeleton of the shape. Indeed, the noise present in Figure 4.2 is still there in Figure 4.4 (left column, middle) and the diagonal lines are still thick. To remedy the situation, the output of Algorithm 3 may be further thinned to obtain an even better approximation, one that will contain lattice points no further than $\sigma/2$ from $\mathrm{Sk}(X)$ everywhere, except possibly near end-points and junction points.

The idea is simple: identify discrete end-points, disallow their removal by anchoring them (similarly to Vincent [54]), and perform a homotopy thinning on $S_{Thick(d)}(X)$; see Algorithm 4. However, devising criteria for the selection of end-points in $S_{Thick(d)}(X)$ turns out to be non-trivial. In the following section we shall examine this problem and propose a solution.

4.2.4 Selection of End-Points

The skeleton approximation $S_{Thick(d)}(X)$ obtained in the previous subsection is "thick" while the desired approximation should not be. The solution is to thin it further by designating certain points as anchors for which the natural choice are the discrete end-points; Algorithm 4 contains the precise idea. To complete the procedure, it then suffices to define the criteria for end-point selection. This will be done by exploiting the defining characteristic of a continuous end-point but, as

Start with $U = S_{Thick(d)}(X) \subset \mathbb{L}_\sigma$
while There is a removable $P \in \partial U$ **do**
  Choose $P \in \partial U$ closest to ∂X among remaining choices.
  if P is not an end-point **and** $NG(P)$ is a tree **then**
    $U = U - \{P\}$
  end if
end while
Return $S_{Thin(d)}(X)$

ALGORITHM 4. Thin skeleton algorithm—it needs the output of Algorithm 3. $NG(P)$ denotes the neighborhood graph of P as described in Section 4.2.2.

we shall see, the range of acceptable values for Thr (d) will be limited as a consequence.

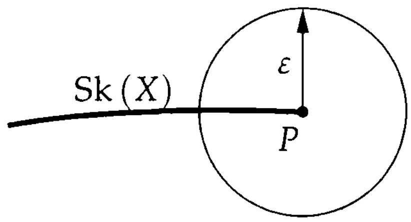

FIGURE 4.5. A skeletal end-point in $\mathbb{R}^2$.

Consider a skeletal end-point, e.g. as depicted on Figure 4.5, and recall Definition 2.20(ii). The statement of this definition expresses the fact that one can place a small enough circle on the end-point which intersects the skeletal curve exactly once. Analogously, on the discrete lattice $\mathbb{L}_\sigma$, one may define as end-points all those $P \in \mathbb{L}_\sigma$ for which there is a small enough discrete circle centered at P and which intersects the (thick) skeleton exactly once. The problem, however, is that of making precise the notion of "small enough" in this context.

By its nature, $L_\sigma(X) \subset \mathbb{L}_\sigma$ only admits a finite number of discrete circles at any given $P \in L_\sigma(X)$ and the one with smallest nontrivial radius r is when $r = \sigma$. However, the above criterion using $r = \sigma$ on $L_\sigma(X)$ is satisfied by many points which should not be labeled as end-points; for example, every corner on the thick skeletons on Figure 4.6 LOWER RIGHT. On the other hand, if r is bigger, then certain end-points may be ignored and whole branches lost; see Figure 4.6 LOWER LEFT. Fortunately, if one is willing to impose conditions on d (and limit the types of skeletons produced as a result), the notion of a discrete end-point may be made precise

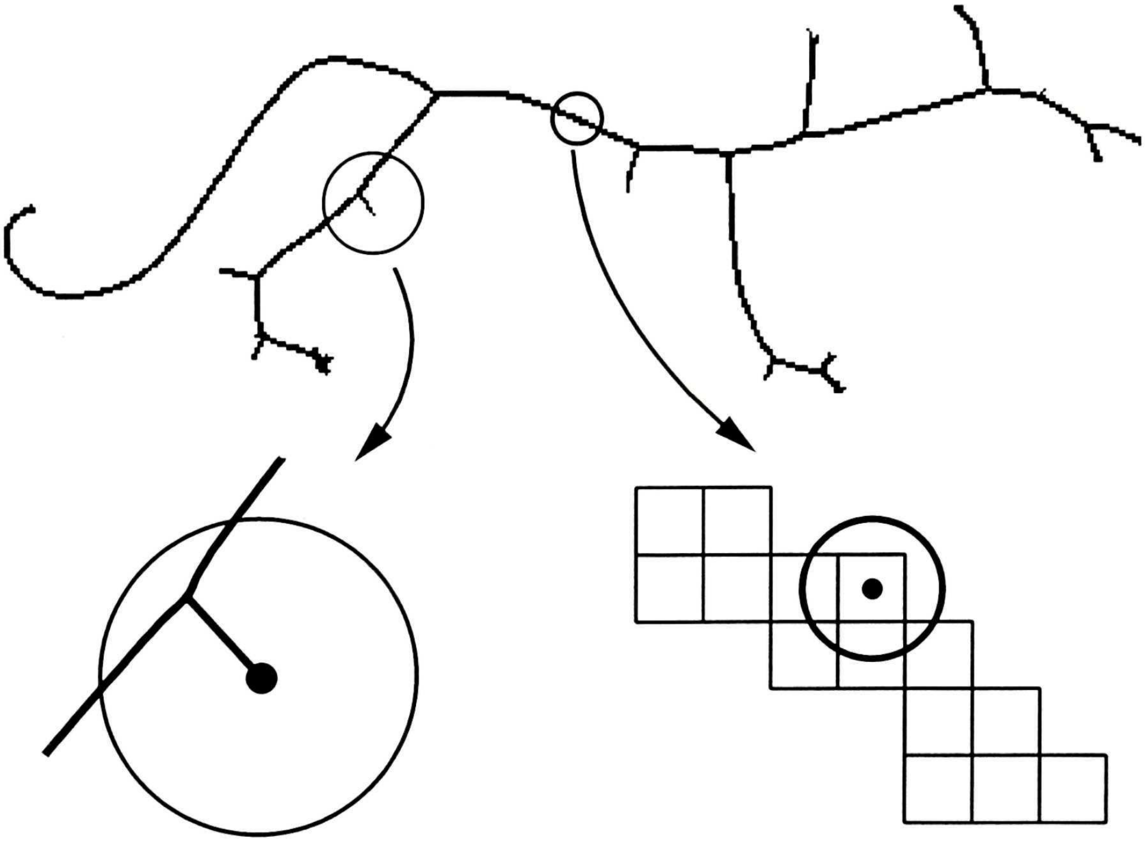

FIGURE 4.6. $S_{Thick(5)}(X)$. Problems with selecting discrete end-points.
LOWER LEFT: The end-point may not be detected if r is too big.
LOWER RIGHT: If r is small some points may be misclassified as end-points.

3.4 Circular Neighborhoods 34
3.4.1 Regular Skeletal Points 34
3.4.2 Skeletal End-Points 37
3.4.3 Skeletal Junction and Pseudo-Junction Points 38
3.4.4 Non-Skeletal Points 41
3.5 Summary 41

4 Flux-Based Skeletonization **43**
4.1 Computational Setup 44
4.1.1 ∂X, ∇D_X and Discrete AOF 44
4.1.2 The Discrete Lattice $\mathbb{L}_\sigma$ 46
4.2 The Algorithm 47
4.2.1 Thresholding the AOF on $\mathbb{L}_\sigma$ 48
4.2.2 Homotopy Thinning 52
4.2.3 Thick Skeletons 55
4.2.4 Selection of End-Points 57
4.2.5 The Final Product: Skeletal Branches 60
4.3 Thresholding on End-Points 62
4.4 Shifting 64
4.5 Experimental Results 66
4.6 Summary 68

Conclusion **74**

A Proofs for Chapter 2 **77**

B Proofs for Chapter 4 **84**

and the aforementioned issues circumvented.

Suppose $d \leq \sigma$ and consider $S_{Thick(d)}(X)$. Then, by the Threshold Lemma (on page 51), $S_{Thick(d)}(X)$ is no more than 2σ thick; in particular, near end-points, there may be no more than two adjacent lattice points in $S_{Thick(d)}(X)$. Hence, any potential discrete end-point P is in one of the following situations

(a) P has a single neighbor or two neighbors which are 4-adjacent to one another,

(b) P has a neighbor Q such that with Q removed, P is as in (a).

The only potential ambiguity in this characterization is when (b) holds for both P and Q. In that case, it is reasonable to pick arbitrarily between the two. Thus, a complete criterion is available that does not mislabel discrete points, as guaranteed by (a), and does not allow for Algorithm 4 to shorten the approximation (e.g. by removing a branch) because of (b). An example of the complete procedure comparing it to Algorithm 3 is shown in Figure 4.7.

4.2.5 The Final Product: Skeletal Branches

A very useful approximation of the skeleton of a shape consists of individual (thin) approximations of the skeletal curves. Once $S_{Thin(d)}(X)$ is obtained, it is easy to "break" it into skeletal branches using simple criteria for labeling discrete points as one of three types: end-point, branch point (analogous to regular points in $\mathrm{Sk}(X)$) and junction points. It is enough to consider the number of intersections of a discrete 1-circle, i.e. the boundary of a 3x3 square, with $S_{Thin(d)}(X)$:

- $P \in S_{Thick(d)}(X)$ is a discrete *end-point* if and only if the square intersects $S_{Thick(d)}(X)$ exactly once;
- P is a *branch point* if there are exactly two intersections; and
- P is a *junction point* otherwise.

Notice that this characterization is naturally broken into discrete curves: by definition, the collection of neighboring branch points gives such a partitioning. Thus, a skeletal curve is approximated by the set of 8-connected branch points and, if so desired, by the two end-points or junction points neighboring this set. Therefore, tangents to the skeleton may be approximated. Observe, however, that

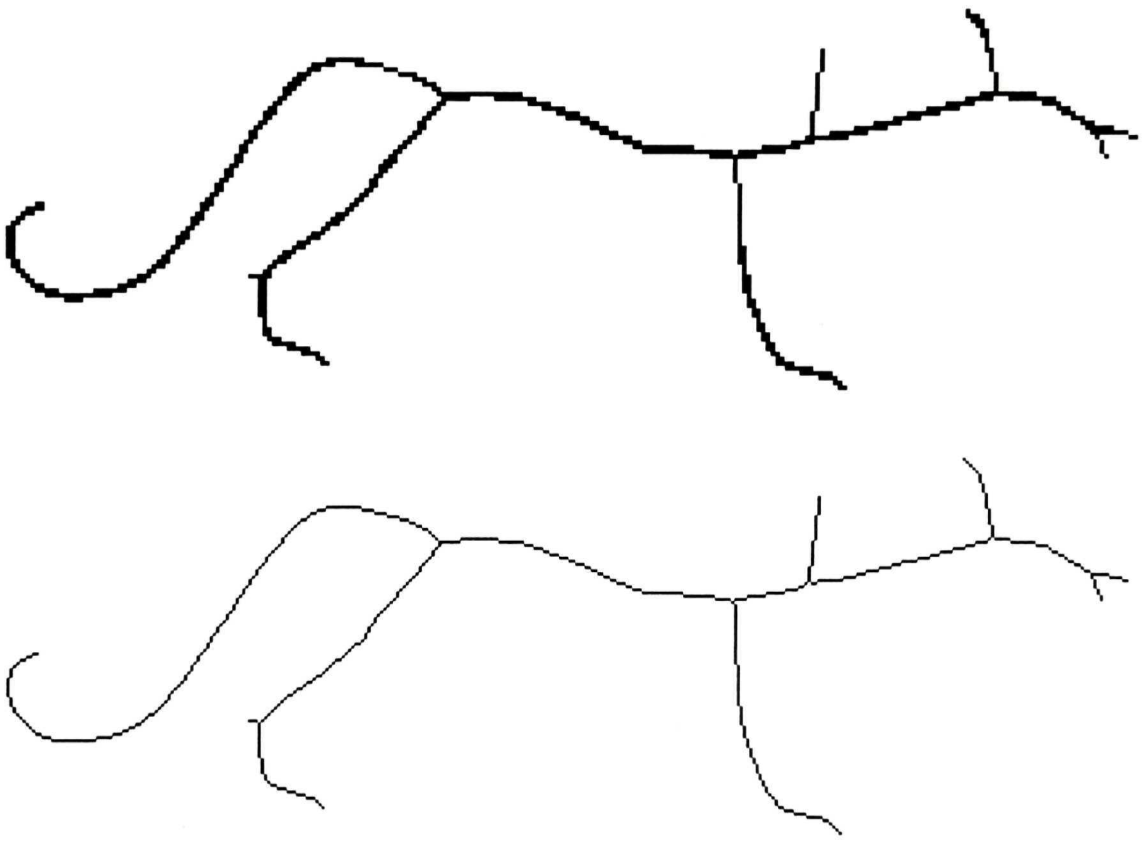

FIGURE 4.7. TOP: Thick skeleton $S_{Thick(1)}(X)$. BOTTOM: The thin discrete skeleton $S_{Thin(1)}(X)$.

it is possible to get several neighboring junction points, but there may be no more than four arranged in a square.

An important question about this representation is: what does it describe? The next section provides an answer.

4.3 Thresholding on End-Points: Consequences for Shape Representation

In Section 4.2.4 a compromise had to be made when determining which discrete points should be labeled as end-points, namely, the threshold Thr (d) had to be taken so that $d \leq \sigma$. Even though at first glance this may seem overly restrictive, it turns out that $S_{Thin(d)}(X)$ (for $d \leq \sigma$) represents a significant portion of the shape X. To demonstrate and quantify this claim, $S_{Thin(d)}(X)$ will be treated as approximating an α-skeleton of X (refer to Definition 2.30 on page 26) and upper bounds for this α will be determined.

To begin, observe that the end-points in $S_{Thin(d)}(X)$ are the first ones to satisfy the threshold as guaranteed by the criterion in Section 4.2.4. Hence, this approximation of the skeleton completely covers the α-skeleton where α is the largest object angle attributable to a discrete end-point in $S_{Thin(d)}(X)$. It is straightforward to obtain an upper bound for such an α using Equation 4.4 of the Threshold Lemma (page 51), reproduced here for convenience:

$$\arcsin\left[\frac{1}{4}\left(2\pi\mathrm{Thr}\,(d) - E\right)\right] \leq \alpha \leq \arcsin\left[\frac{1}{2\sqrt{3}}\left(2\pi\mathrm{Thr}\,(d) + E\right)\right]$$

where $E = 4\delta + \frac{4}{3}\frac{\pi^3}{n^2}$ and the object angle is in the interval $[\alpha - \delta/2, \alpha + \delta/2]$. With $E = 0$ and $d = \sigma$, the object angle is bounded above by

$$\alpha \leq \arcsin\left(\frac{2\pi\mathrm{Thr}\,(\sigma)}{2\sqrt{3}}\right) \approx 36.226^\circ$$

and the α-Skeleton Theorem (see page 27) implies that this α-skeleton represents over 96% of the shape's area.[4] The question now is: does $S_{Thin(d)}(X)$ approximate well the α-skeleton for $\alpha = 36.226^\circ$?

[4] Actually, the theorem gives this bound for individual protrusions; thus, this it is a loose bound for most shapes.

The answer is yes if the the end-points of the α-skeleton are at least distance b from the boundary and if σ is appropriately chosen. Then, as we shall now see, it can be shown that the discrete point will be at most σ from the skeletal location which has object angle greater than this α.

To see this, let $P \in \mathrm{Sk}\,(X)$ be within the circle of radius σ around a discrete end-point in $S_{Thin(\sigma)}(X)$ (such a P must exist as guaranteed by the Threshold Lemma). Denote by Γ_P the set of all shortest paths connecting P to a real end-point of $\mathrm{Sk}\,(X)$ and pick some $\gamma \in \Gamma_P$; see Figure 4.8 TOP. First, suppose any circle of radius σ containing portions of γ actually contains line segments. Now, consider walking on γ starting at the real end-point and moving toward P. If, along the way, the object angle becomes greater than α and then decreases, denote by $Q \in \gamma$ the location where it is greatest. Assume, without loss of generality, that $\mathrm{d}\,(P, Q) > \sigma$. The greatest rate of decrease of the object angle on the stretch starting at Q and ending at P, denoted γ_{QP}, is reached if all $P' \in \gamma_{QP}$ are closest to the same boundary point B. This rate is given by $\alpha'(s) = \frac{\sin\alpha(s)}{r(s)}$ where $\alpha(s)$ and $r(s)$ are the object angle and the distance to the boundary, respectively; see Figure 4.8 BOTTOM.

Hence, on the circle centered at P and of radius σ, δ is no more than if the circle was in a gradient field defined by a single point at least b away from its center. By Lemma 4.3, it follows that δ is bounded as

$$\delta \leq 2\arcsin\left(\frac{\sigma}{\sigma+b}\right)$$

Now, if γ is not a line segment, we can use the fact that the skeleton's curvature is no more than one over the radius value there (see Damon [15]) and choose the lattice spacing, so that any portion of γ contained within a circle of radius σ has tangents with maximum variation in orientation ζ for any specified ζ.[5] So, in the general case,

$$\delta \leq 2\arcsin\left(\frac{\sigma}{\sigma+b}\right) + \zeta\,.$$

Hence, given δ_0 and using the above, σ may be chosen so that $\delta \leq \delta_0$. Moreover, Lemma 4.5 suggests a method for bounding the numerical error, i.e. how to choose n in order to guarantee $\mathrm{NumErr}(n) \leq \varepsilon \in \mathbb{R}_+$. Consequently, given $E_0 > 0$, we

[5]For example, if the boundary is given by a sequence of evenly points, say on a grid, then the curvature of the skeleton may not be more than one over that separation.

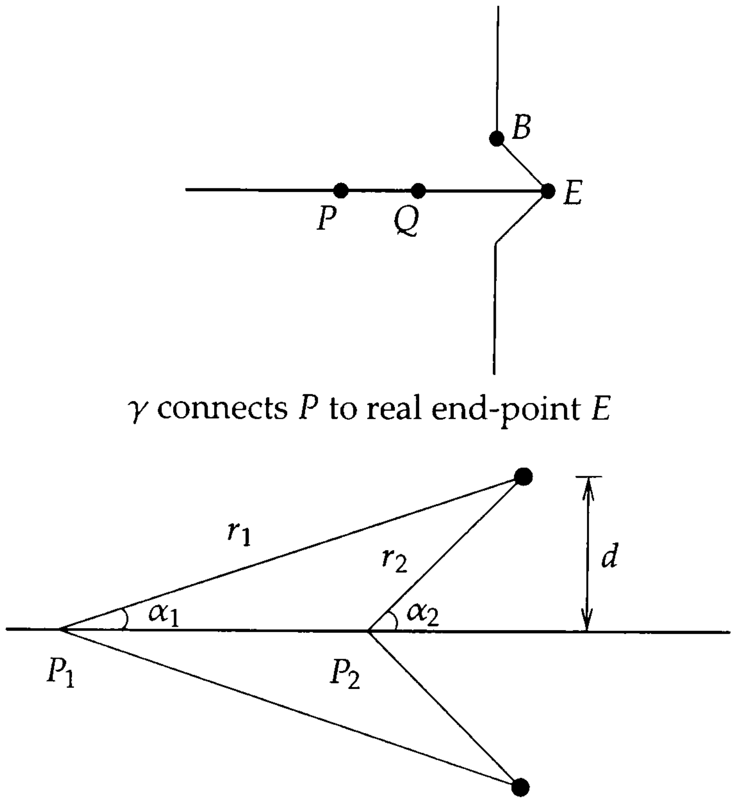

FIGURE 4.8. TOP: boundary (thin line) and corresponding skeletal curve (thick line). BOTTOM: The object angle, denoted by α_i for $i = 1, 2$, can be expressed as ratios and the derivative of the object angle can be obtained by looking at the difference $\alpha_2 - \alpha_1$ and taking the limit as $d(P_1, P_2) \to 0$. Thus, $\alpha' = \sin(\alpha)/r$.

can make $E \leq E_0$ and

$$\frac{1}{2\pi}4\sin(\alpha_P) + E_0 \leq \left|\mathscr{F}_\sigma^{(n)}(P)\right| \leq \frac{1}{2\pi}(4\sin(\alpha_P) - E_0)$$

where $\left|\alpha_P - \alpha_Q\right| \leq \frac{\sin\alpha_Q}{b}\sigma$.

4.4 Shifting

Under certain conditions, the discrete approximation $S_{Thin(d)}(X)$ of the skeleton may be further improved. In fact, it is possible to shift the discrete points to be arbitrarily close by applying a bisection scheme, see Algorithm 5. A possible side effect of this algorithm is a reordering of the discrete points as obtained by the procedure outlined in Section 4.2.5. In general, it may be very difficult (impossible)

INPUT: The tolerance τ and $S_{Thin(d)}(X) \subset \mathbb{L}_\sigma$.
for all $P \in S_{Thin(d)}(X)$ **do**
  $\mathbf{v}_P = \nabla D_X(P)$
  $s = \sigma$
  while $\nabla D_X(P + s\mathbf{v}_P) = \mathbf{v}_P$ **do**
    $s = s + \sigma$
  end while
  $Q = P + s\mathbf{v}_P$
  while $\mathrm{d}(P, Q) > \tau$ **do**
    $M = \frac{1}{2}(P + Q)$
    if $\nabla D_X(M) \neq \mathbf{v}_P$ **then**
      $Q = M$
    else
      $P = M$
    end if
  end while
end for
Return $S_\tau(X)$

ALGORITHM 5. Shifting of the discrete skeleton approximation given by $S_{Thin(d)}(X)$, the output of Algorithm 4 (see page 58). Returns a discrete skeleton approximation where the points are no more than τ away from $\mathrm{Sk}(X)$.

to put them back in order, but if the smallest object angle in the α-skeleton is not too small relative to the lattice spacing, Algorithm 6 will do just that.

The property this algorithm relies on is that the closest point Q to a given P in the original unordered set will, in fact, be the next point along the skeletal curve. This is true whenever $\mathrm{d}(P, Q)$ is small relative to the curvature of the curve. So, if $\mathrm{d}(P, Q)$ is always less than or equal to $\mathrm{d}(P, \partial X)$, the procedure will yield an appropriately ordered set O.

Now, let's derive a criterion that ensures this and allows us to estimate $\mathrm{d}(P, Q)$. Observe that, before shifting, $\mathrm{d}(P, Q) \leq \sqrt{2}/2\sigma$ and that a discrete point will not be shifted more than $\frac{\sigma}{2\sin\alpha}$ from its original position where α is the angle of the gradient through the point and the skeleton. Therefore, if the smallest object angle is α_0, choosing σ such that

$$\frac{\sigma}{2\sin\alpha_0} < \frac{1}{2}\frac{\sqrt{2}}{2}\sigma$$

ensures that Algorithm 6 will produce ordered sequences of points. Notice that a

INPUT: The shifted thin skeleton $S_{\mathcal{T}}(X)$.
OUTPUT: Set of ordered curves $\mathcal{O}$.
for all Curves $\mathcal{C}$ in $S_{\mathcal{T}}(X)$ **do**
  Pick $P \in \mathcal{C}$
  Pick $Q \in \mathcal{C} - \{P\}$ closest to P.
  Pick $Q' \in \mathcal{C} - \{P, Q\}$ closest to P.
  if $\mathrm{d}\,(Q,\ Q') < \mathrm{d}\,(P,\ Q')$ **then**
    $\mathcal{O} = \{P, Q, Q'\}$
  else
    $\mathcal{O} = \{Q', P, Q\}$
  end if
  while $\mathcal{C} - \mathcal{O} \neq \varnothing$ **do**
    Let P be the beginning of $\mathcal{O}$ and Q its end.
    Pick $P' \in \mathcal{C} - (\mathcal{O} \cup \{P\})$ closest to P.
    Pick $Q' \in \mathcal{C} - (\mathcal{O} \cup \{Q\})$ closest to Q.
    if $\mathrm{d}\,(Q,\ Q') < \mathrm{d}\,(P,\ P')$ **then**
      insert Q' in $\mathcal{O}$ after Q
    else
      insert P' in $\mathcal{O}$ before P
    end if
  end while
end for
Return $S_{\mathcal{T}}(X)$

ALGORITHM 6. Reordering of points after shifting; see Algorithm 5.

lower bound for α_0 may be obtained by assuming that the closest distance between any two dull corners is no more than d_c, then

$$\frac{d_c}{R} \leq \sin(\alpha_0)$$

where R is the radius of the shape X, i.e. such that X is contained in a ball of radius R. This bound is particularly useful when the initial shape is given on a discrete grid; then, d_c is no more than one grid separation.

4.5 Experimental Results

In this section, we shall perform experiments using Algorithm 4 and Algorithm 5. We shall apply the boundary reconstruction procedure suggested in Section 2.4

on the output of a thin, shifted discrete approximation of the skeleton in order to, qualitatively, corroborate the theoretical guarantees provided in this chapter.

The setup is the same as in Section 4.1 where the shapes are given by boundaries consisting of line segments. The samples do not have holes in them, but that is not a limitation of the algorithms, it is a consequence of the current implementation. Figure 4.9 presents an overview of the steps involved in the experiments. First, a binary image is taken and its discrete contour traced. This gives an ordered list of pixels connected with line segments. Thus, a closed curve $C(t) : \mathbb{R} \to \mathbb{R}^2$ is obtained[6]. Next, Algorithm 4 is applied ($d = 1$), which yields a one-pixel thick discrete skeleton. These are shifted using Algorithm 5 to better approximate the tangents to $\mathrm{Sk}(X)$. Finally, the object angle is obtained through the average outward flux approximation as predicted by Equation 3.7 (see page 37) and, for each discrete skeletal point more than two pixels away from a junction point or an end-point, the two corresponding boundary points are reconstructed using Equation 2.5 (see page 25).

Figure 4.10 provides additional examples of the above computation and Figure 4.11 compares the accuracy of the method to that of an exact calculation. The profile on the left uses straight lines to show the association of regular skeletal points with their bi-tangent points on the contour (the black circles). Here the association has been determined by using the average outward flux limit values to obtain the object angle. The profile on the right demonstrates an "exact" computation, where the bi-tangent points are obtained by connecting each regular skeletal point from the shifted approximation to its two closest contour points. Notice how similar the two computations are.

The final demonstration is shown on Figure 4.12. A rectangle defined on the discrete grid is rotated on the lattice and the above experiment run for three angles. Observe that, for all samples, the accuracy of the reconstruction suggests an even higher accuracy for the skeleton approximation. The reason for this is that both tangents to the skeleton and object angle have to be estimated well in order to obtain a good approximation to the boundary points: a small inaccuracy in either one results in a big error for boundary points, especially where the radius function is highest.

It should be pointed out that certain portions of the contour have not been reconstructed because the end-points of the skeleton have been chosen to satisfy an

[6]The pixel locations are smoothed to account for jaggedness inherent in all discrete images.

object angle threshold (above 30°). Consequently, the end-points shown, although very close to the real ones, may not be actual end-points of the skeleton and whole branches of Sk (X) may not have corresponding discrete branches.[7] Thus, to approximate the missing portions of the contour, it would be necessary to draw the circular arcs corresponding not only to the approximate end-points, but to account for the missing branches as well. The latter being a nontrivial task, the reconstruction was limited to the regular skeletal points.

4.6 Summary

In this chapter, we developed a method that computes an approximation to the α-skeleton of a shape and demonstrated how to choose parameters in order to get an arbitrarily good representation of the continuous object. The main idea revolves around adapting the AOF criterion discussed in Chapter 3 to the discrete lattice. We analyzed the effects of thresholding the discrete version of the AOF and determined a threshold function that ensures only points no more than aa specified distance from the skeleton may be kept. Then, we incorporated this idea into a homotopy thinning procedure which outputs a thick but topologically adequate approximation. Once we identified the discrete end-points, we could thin again and obtain a thin discrete skeleton which can be naturally broken into discrete skeletal branches. These branches are simply an ordered sequence of points, hence, tangents to the medial axis may be obtained from them. To approximate those tangents even better, we then devised a shifting procedure which moves the points in a branch to be within a specified (arbitrary but strictly positive) distance to the skeleton while ensuring that the ordering is preserved. Algorithm 7 summarizes these steps.

[7]For example, on Figure 4.4, the panthers paws have small branches when thresholding with $d = 3\sigma$ which are not detected by the lower threshold of $d = \sigma$.

INPUT: $L_\sigma(X) \subset \mathbb{L}_\sigma$
OUTPUT: $S_{\mathrm{Thr}(d)}(X) \subset \mathbb{L}_\sigma$
Get $S_{Thick(1)}(X)$ by running Algorithm 3 on $L_\sigma(X)$
Get $S_{Thin(1)}(X)$ by running Algorithm 4 on $S_{Thick(1)}(X)$
Shift $S_{Thin(1)}(X)$ within τ of $\mathrm{Sk}(X)$ with Algorithm 5
Ensure branches are ordered with Algorithm 6 and output $\mathrm{Sk}_\sigma(X, 37^\circ)$

ALGORITHM 7. The complete procedure.

List of Figures

1.1 TOP: The grassfire process. BOTTOM: The result. 4
1.2 3D surface under image plane induced by the distance function. . . . 8

2.1 Disallowed shape 14
2.2 Necessity of condition (i) in Definition 2.6. 19
2.3 An examples of a sharp corner (LEFT), and a dull corner (RIGHT) . . . 20
2.4 Examples of the different types of skeletal points. 21
2.5 Object angle 24
2.6 The proportion of area represented by an α-skeleton 27

3.1 A region R which intersects a branch of the skeleton S. 31
3.2 The object angle $\alpha = \alpha(P)$ at a regular skeletal point P. 35
3.3 AOF simplified. 35
3.4 A circular neighborhood of radius ε around the end-point P. 39
3.5 Junction point types. 40

4.1 Example of $L_\sigma(X)$. 45
4.2 Skeletonization by simple thresholding 53
4.3 3x3 neighborhood 53
4.4 Thick skeletons 56
4.5 A skeletal end-point in $\mathbb{R}^2$. 58
4.6 End-point selection problems 59
4.7 Thin skeleton. 61
4.8 Consequences of thresholding on the AOF 64
4.9 Experiments: steps 70
4.10 Experiments: results 71
4.11 Experiments: comparison 72
4.12 Examples: stability under rotation 73

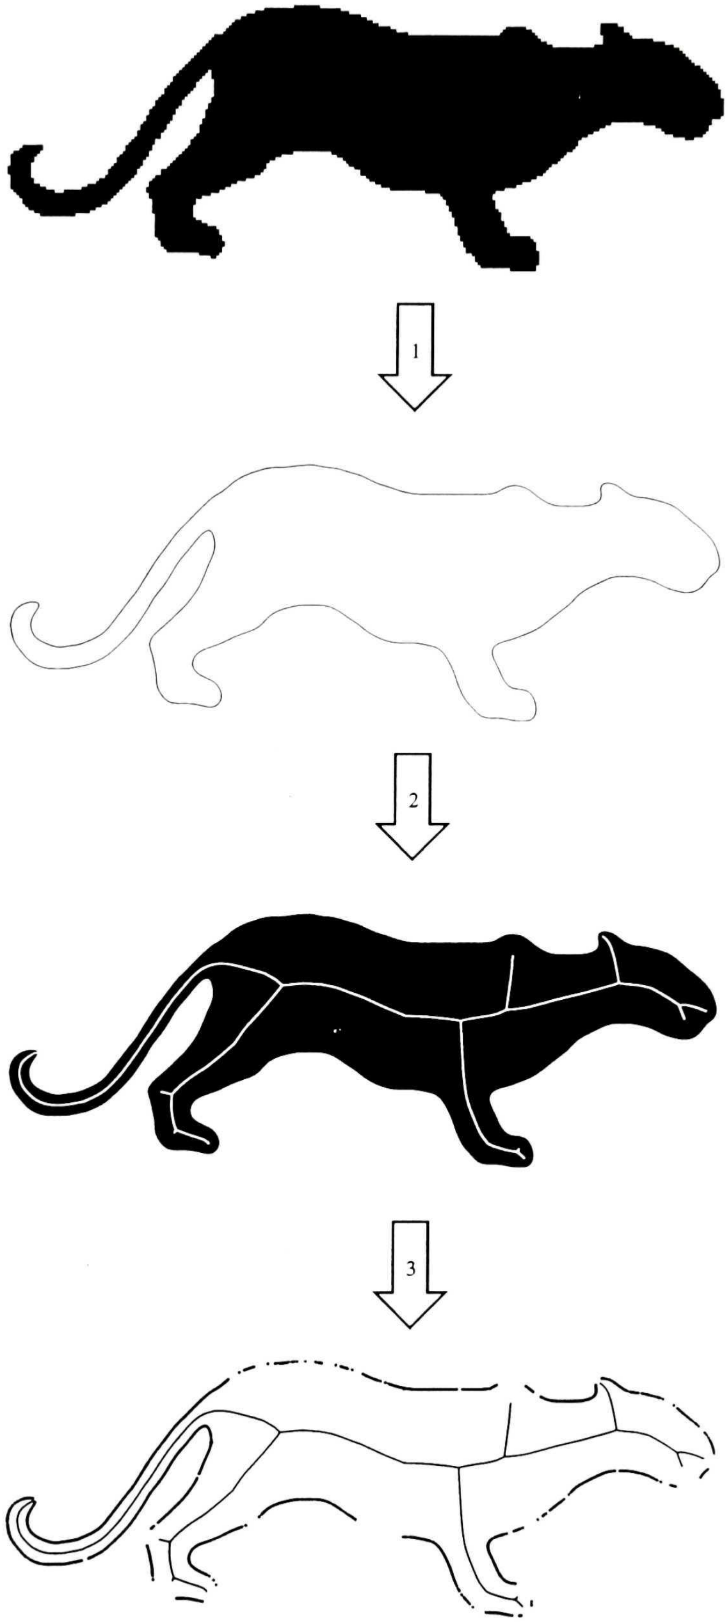

FIGURE 4.9. (1) From a binary image, the boundary is extracted and represented as a continuous curve. (2) The skeleton is computed and shifting performed. (3) Using the average outward flux and radius values along the skeleton, the boundary is reconstructed.

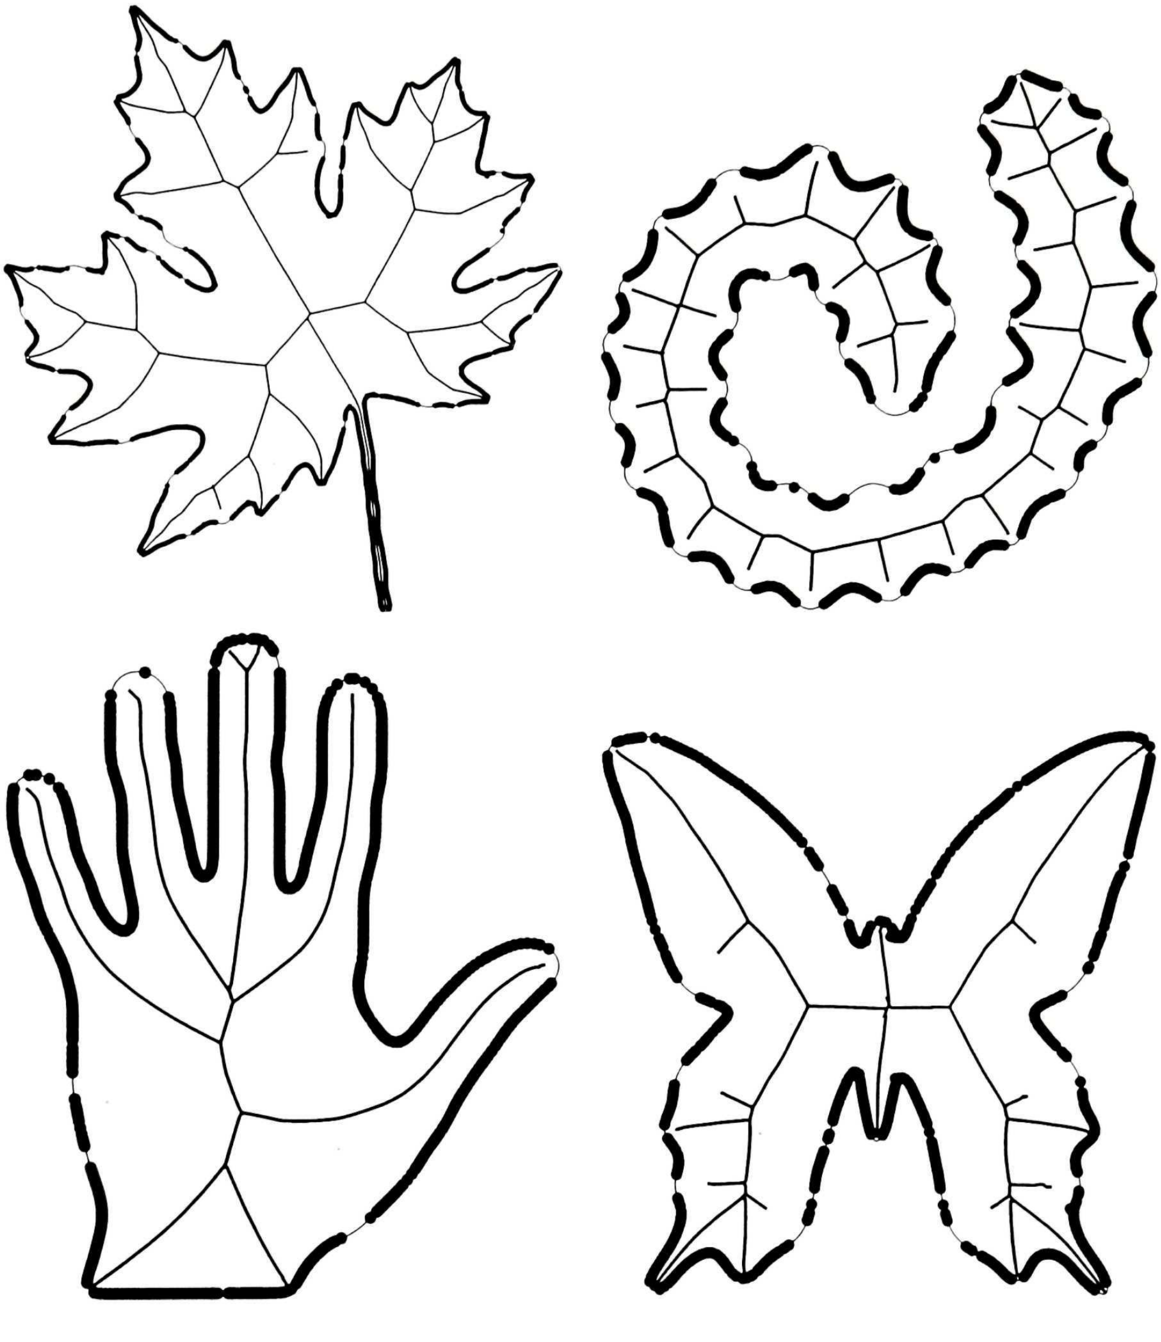

FIGURE 4.10. For each shape the original boundary is shown as thin curve, the skeleton obtained using the average outward flux is shown with thick curves and the boundary points estimated from the skeleton using the relationship between the average outward flux and the object angle α are shown with black circles.

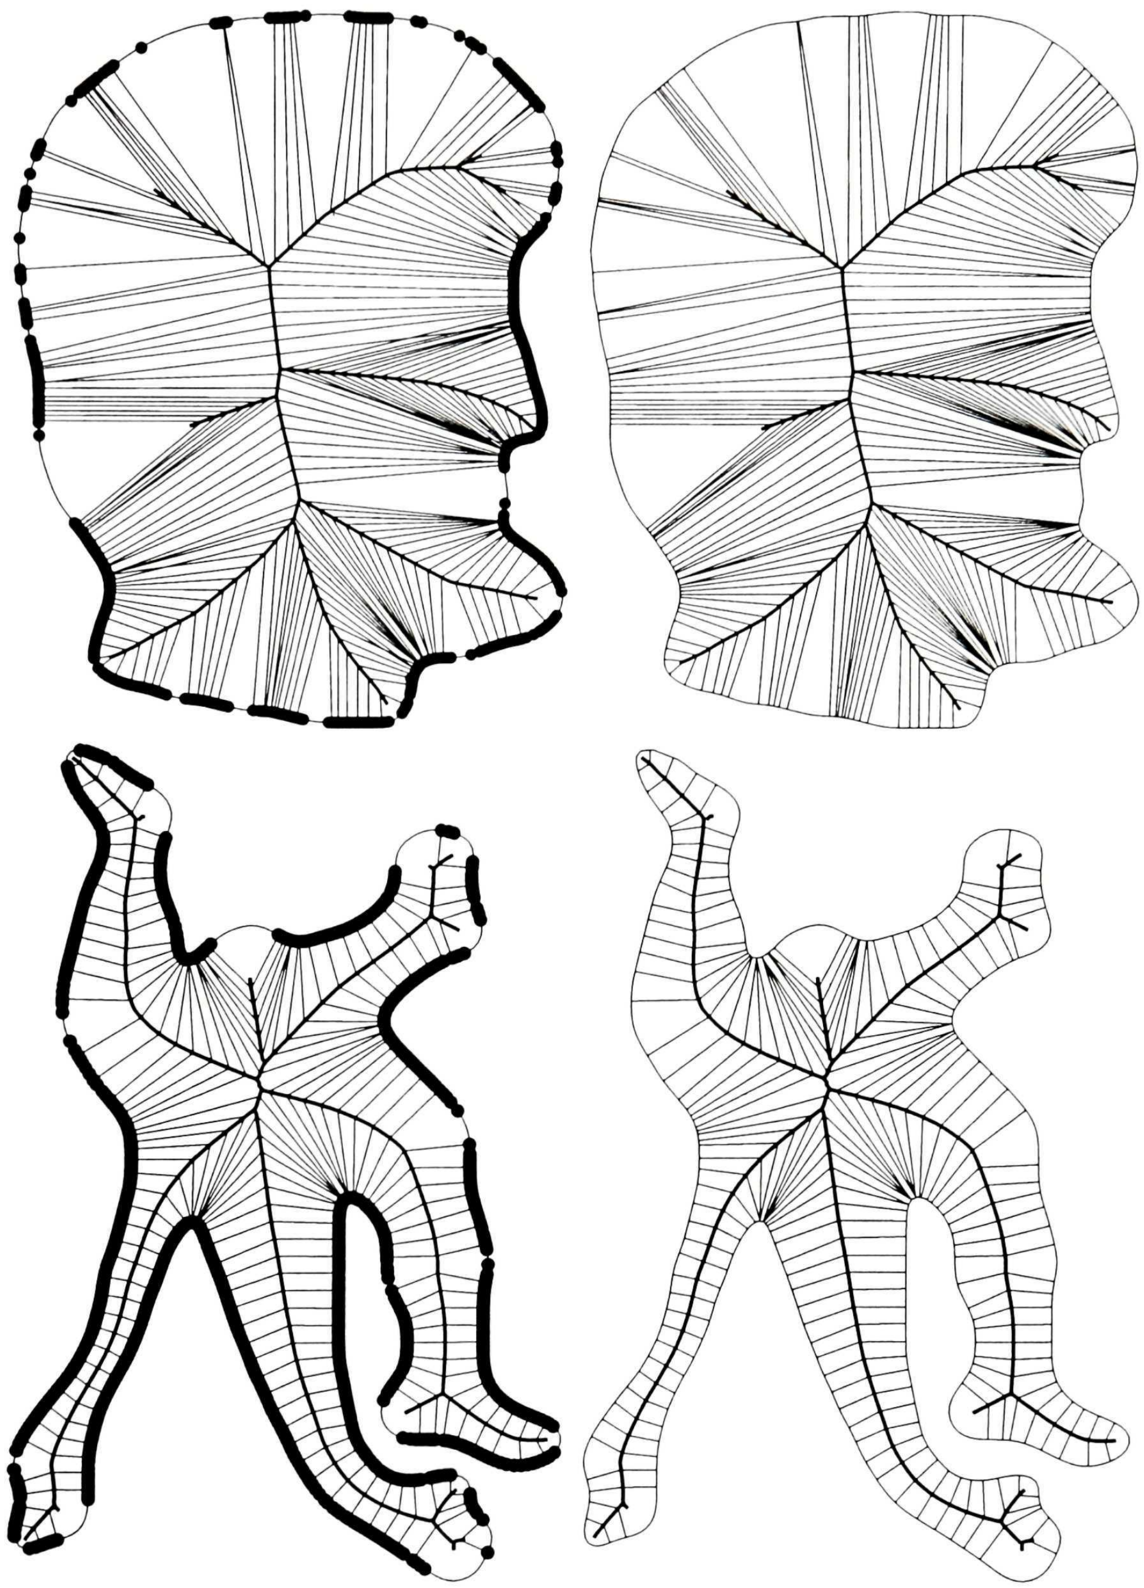

FIGURE 4.11. Bi-tangent points associated with regular skeletal points: (LEFT) computed using average outward flux information and (RIGHT) computed explicitly.

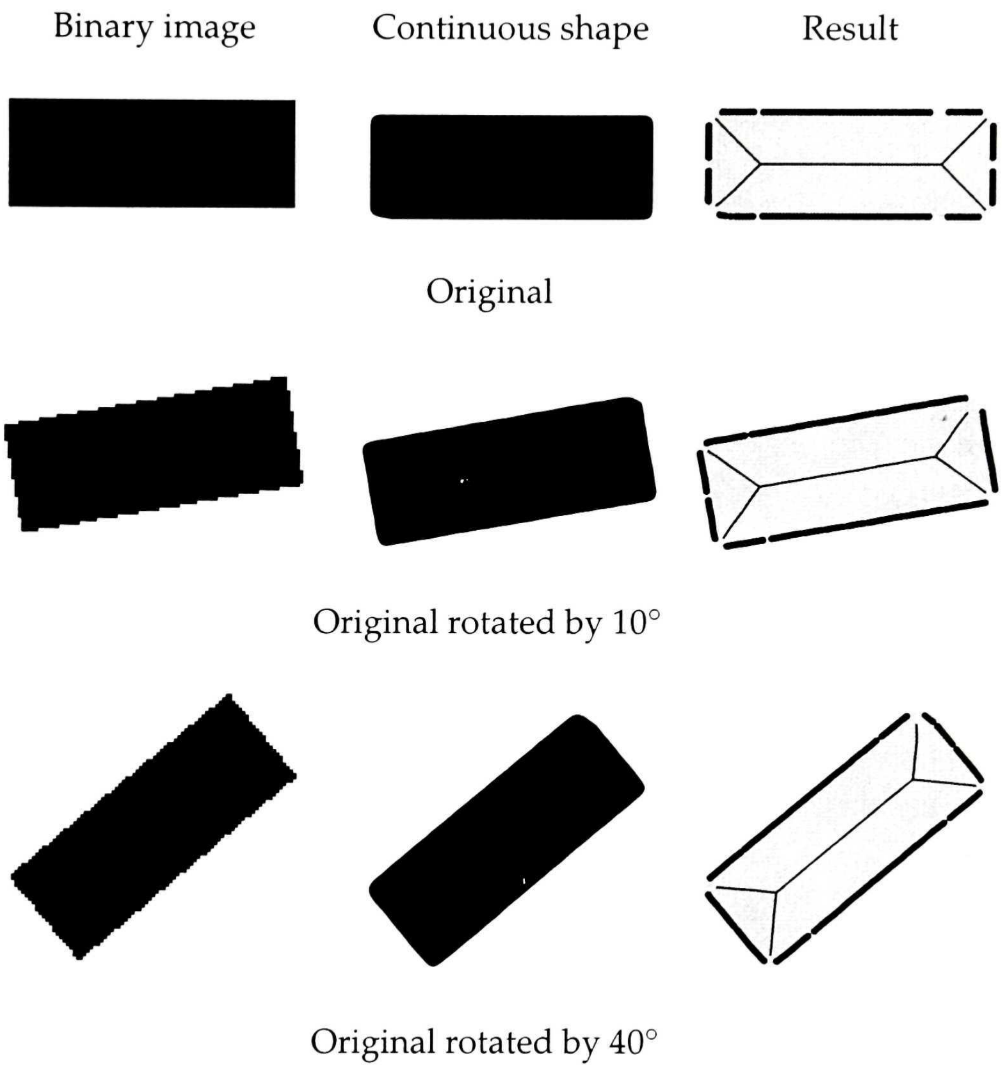

FIGURE 4.12. Examples showing the boundary reconstruction of a shape that has been rotated on the discrete grid.

Conclusion

Overview

In this thesis we have studied the problem of computing the medial axis of a 2D shape. The discussion began in Chapter 2 where the necessary formalisms were introduced and several properties of skeletons demonstrated. Then, in Chapter 3, we saw how a normalized flux measure could be used to identify medial points distinguishing them from non-medial points. In the previous chapter, Chapter 4, we adapted those results to the discrete lattice $\mathbb{L}_\sigma$ and developed an algorithm for estimating the skeleton by a finite number of "sample" points. Here, we shall revisit each of these steps.

The first chapter in the development of this thesis, Chapter 2, introduced a number of definitions and properties about shapes and their skeletons. As is necessary for the implementation on digital machines, we adopted a formalism for shape that allows its skeleton to be seen as a finite graph. In turn, the medial axis was introduced in two different ways—as the locus of maximally inscribed disks and as singularities of D_X—both of which provided a unique perspective of the same mathematical object. We have since mixed language, referring to the "radius" function along skeletal branches as borrowed from the first definition and to object "angles" inspired by the second one. The interplay between these formulations also allowed us to establish a number of useful properties of the skeleton: it is thin (trivial interior); with finitely many non-regular points and, consequently, finitely many smooth curves; it has a natural interpretation as an invariant under rigid transformations; and it can be used to represent the boundary explicitly, without having to take the envelope of maximal disks (see Theorem 2.28). We were then able to talk about shapes and their skeletons formally having developed most of the necessary language.

In Chapter 3, we set out to derive a criterion which identifies points on the plane as either belonging to the skeleton or not. There, the second definition of medial axis—singularities of D_X—proved most useful. Indeed, the main idea was to locate the points in $\mathbb{R}^2$ where the distance function is not differentiable. To do this, we demonstrated that the flux through the boundary of an appropriately shrinking region behaved (asymptotically) differently if the region collapsed to a point on the skeleton than to a point away from it. In fact, we showed that the average outward flux, obtained by dividing the flux by the length of the boundary, exhibited the behavior explicitly: it is non-zero only on medial points. We demonstrated that this value was actually the sine of the object angle for general shrinking regions, and we derived the explicit behavior for circular regions on all types of skeletal points (see Table 3.1 on page 42). However, these results could not be used without modification on a discrete lattice.

In the final chapter of the development (Chapter 4), we adapted the criterion obtained earlier for use on the discrete lattice and then applied it to obtain an algorithm which computes an arbitrarily close approximation to the skeleton. There were several steps. First (Section 4.2), we analyzed the behavior of the average outward flux through a small circle and determined how to threshold this value so as to define a, no more than two pixel thick, anchor set for a homotopy thinning. That allowed us to extract end-points of the thick approximation and thin again. We then had a one pixel thick approximation (except possibly at junction points) that was guaranteed 7to be no more than half a pixel away from the real skeleton. Further, we showed (in Section 4.3) that this thin approximation actually covered the α-skeleton (for $\alpha = 37°$) evenly which demonstrated the adequacy of this representation. In order to obtain an even closer approximation, (in Section 4.4) we provided the necessary procedures to shift the discrete points obtained previously within an arbitrary (but positive) distance to the skeleton and we showed how to guarantee that the shifted points cover the skeleton appropriately. The effectiveness of this approach was then demonstrated qualitatively in Section 4.5 by reconstructing the boundary using the shifted thin skeleton.

In summary, this thesis makes two main contributions. First, the results presented in Chapter 3 show that the limit of the average outward flux through a region (e.g. circle) shrinking to a point distinguishes between medial points and non-medial points; it is a function of the object angle in the former case and identically zero in the latter. Hence, the limit of the AOF is an invariant under rigid transformations and can be used to identify skeleton points. Second, using an approximation to this criterion on the discrete lattice, we have presented here a skeletonization algorithm that approximates the α-skeleton (for $\alpha = 37°$) arbitrarily well: it returns finitely many points with user-specified minimum spacing and which are no more than a user-specified distance to the real α-skeleton.

Future Directions

Recall that the selection of end-points from the thick skeleton $S_{Thick(d)}(X)$ could only be done reliably when $d = 1$. However, an extension for higher values of d seem possible. Counting intersections of a discrete circle of radius $d+1$ may overlook some branches, but adapting this radius based on the local structure of the $S_{Thick(d)}(X)$ may solve the problem. Then, Algorithm 7 would approximate an α-skeleton for an even lower α.

Another possible way to obtain such an approximation may be to simply extend $\mathrm{Sk}_\sigma\,(X, 37°)$ obtained by the algorithm presented in this thesis. The approach would use $S_{Thick(d)}(X)$ as a mask and only try to extend $\mathrm{Sk}_\sigma\,(X, 37°)$ for points in $\mathcal{S}_\sigma\left(S_{Thick(d)}(X)\right)$. Thus, for each element P in $S_{Thick(d)}(X)$ not already represented in $\mathrm{Sk}_\sigma\,(X, 37°)$, one could study numerically (e.g. through the Nelder-Mead Simplex method in 1D, see [25]) the local maxima of the distance function along the circle C_σ^P.

Finally, an interesting question unrelated to the computational aspects of the method presented here, but having to do with the representational power of its result is the following: How important are α-skeletons perceptually; that is, is there an α_0 such that no α-skeleton with $\alpha < \alpha_0$ contains more perceptual information than the α_0-skeleton?

Appendix A

Proofs for Chapter 2

Corollary 2.10 is restated below as Proposition A.1 and the proof follows.

Proposition A.1 (Corollary 2.10, p. 16). *Let $P \in \mathbb{R}^2 - \partial X$. Assume $|P_C| = 1$ and let $Q \in \partial X$ be the unique point in P_C. If ∂X is C^0 near Q, then $D_X(P)$ is at least C^1 near P.*

Proof. If ∂X is C^0 but not C^1 near Q, then the normal to ∂X is not well-defined at Q. However, since the boundary of the shape is a closed curve, Q must be where two C^k pieces, B_1 and B_2, meet. Denote by $\mathbf{n}_i$ (for $i = 1, 2$) the limit of the normal vector of B_i while approaching Q along B_i. See Figure A.1. These two vectors span a cone which contains all points closest to Q; hence, P must be in it. If P is strictly inside the cone D_X is C^∞, Theorem 2.9(a) applies and Equation 2.2 holds. On the other hand, if P' were strictly outside the cone, then it would be closer to some point on, e.g. $B_1(s)$ than to Q, assuming w.l.o.g. that P' is sufficiently close to the boundary of the cone. Therefore, Theorem 2.9(b) applies to such P' which shows that approaching the boundary of the cone from outside or from inside yields the same value for $\nabla D_X(P)$. Hence, D_X is C^1 even if P is on the boundary of the cone.

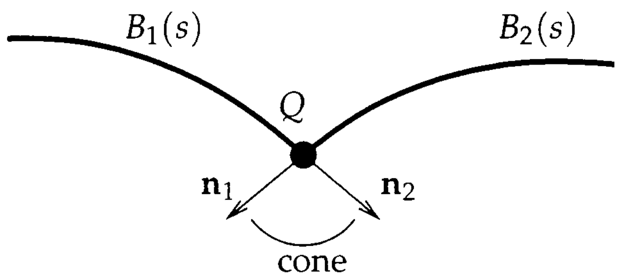

FIGURE A.1. A region R which intersects a branch of the skeleton S.

□

Theorem 2.14 is restated below as Theorem A.2.

Theorem A.2 (Theorem 2.14, p. 18). *Let*

$$E = \{P \in X \; : \; P \text{ is a centre of curvature for } \partial X \text{ and } |P_C| = 1\}.$$

Then,

$$\mathbf{MA}(X) - E = \mathrm{Sk}(X).$$

Proof. If $P \in \mathbf{MA}(X)$ and $|P_C| = 1$ (P_C must contain at least one point), then there exists $r > 0$ such that $\mathcal{B}_r(P)$ is maximally inscribed in X but $\partial\mathcal{B}_r(P)$ and ∂X intersect at a single point Q. Hence, ∂X must be $\mathcal{C}^k$ near Q — otherwise $r = 0$ — in particular, $\partial\mathcal{B}_r(P)$ is an osculating circle to ∂X at Q and P is center of curvature. Therefore, the only points with trivial boundary support included in the medial axis are those in E. Thus, $\mathbf{MA}(X) - E \subseteq \mathrm{Sk}(X)$.

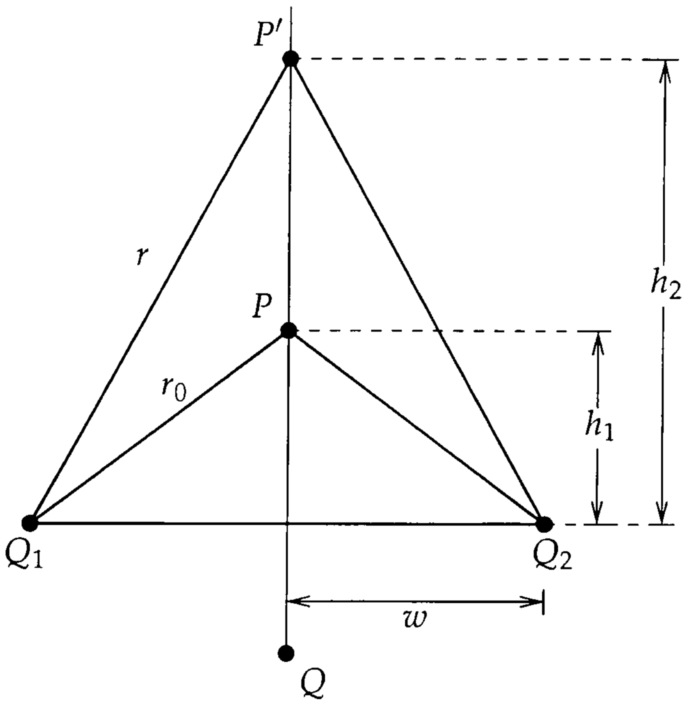

FIGURE A.2. See proof of Theorem 2.14.

On the other hand, if $P \in \mathrm{Sk}(X)$, then, since $\mathrm{d}(P, Q) = D_X(P)$ for any $Q \in P_C$, $\mathcal{B}_{\mathrm{d}(P,Q)}(P) \subseteq X$. Now, if another ball, $\mathcal{B}_r(P')$, contains $\mathcal{B}_{\mathrm{d}(P,Q)}(P)$, then any $Q \in P_C$ must be on the boundary of $\mathcal{B}_r(P')$ because otherwise there is $\varepsilon > 0$ such that $\mathcal{B}_\varepsilon(Q) \subseteq \mathcal{B}_r(P')$; however $\mathcal{B}_\varepsilon(Q) - X \neq \varnothing$ so $\mathcal{B}_r(P') - X \neq \varnothing$. Hence, choosing two distinct points $Q_1, Q_2 \in P_C$ implies $\mathrm{d}(P', Q_1) = \mathrm{d}(P', Q_2) = r$. Thus, two isosceles triangles — $\triangle PQ_1Q_2$ and $\triangle P'Q_1Q_2$ — are defined as in Figure A.2.

Now we show that $r = \mathrm{d}(P, Q_1) = D_X(P)$. Put $r_0 = \mathrm{d}(P, Q_1)$, $\Delta r = r - r_0$, $h = h_2$ and $\Delta h = h - h_1$. Suppose the claim does not hold. Then, $r > r_0$ and $\Delta r, \Delta h > 0$. Since the triangles are isosceles, write $h_1^2 + w^2 = r_0^2$ and $h_2^2 + w^2 = r^2$. Thus,

$$\begin{aligned} h^2 - r^2 = w^2 &= (h - \Delta h)^2 - (r - \Delta r)^2 \\ (h+r)(h-r) &= (h + r + \Delta h + \Delta r)(h - r + \Delta h - \Delta r) \\ \tfrac{h+r}{h+r+\Delta h+\Delta r} &= 1 + \tfrac{\Delta h - \Delta r}{h-r} \end{aligned}$$

Hence, since $h - r < 0$, $\Delta h - \Delta r > 0$ so $\Delta h > \Delta r$. Now, suppose the point Q on Figure A.2 is on the boundary of $\mathcal{B}_{\mathrm{d}(P,\, Q)}(P)$. Then it is $\mathrm{d}(P, Q) = r_0$ and $\mathrm{d}(P', Q) = r_0 + \Delta h > r_0 + \Delta r = r$. Therefore, Q lies outside the boundary of $\mathcal{B}_r(P')$ which contradicts the assumption that $\mathcal{B}_{\mathrm{d}(P,\, Q)}(P) \subseteq \mathcal{B}_r(P')$. This means that $r = \mathrm{d}(P, Q)$ and $P' = P$ which shows that the ball $\mathcal{B}_{\mathrm{d}(P,\, Q)}(P)$ is maximally inscribed and implies that $\mathrm{Sk}(X) \subseteq \mathbf{MA}(X) - E$ completes the proof of the theorem. $\square$

Corollary A.3 (Corollary 2.15, p. 18). *Let X be a shape. The skeleton of X is a collection of bounded curves, i.e.*

$$\mathrm{int}(\mathrm{Sk}(X)) = \varnothing.$$

Proof. It suffices to show that $\mathrm{int}(\mathrm{Sk}(X)) = \varnothing$ since the first statement claims that the skeleton of a shape is either a 1-dimensional object or a set of points. So, suppose that $\mathrm{int}(\mathrm{Sk}(X)) \neq \varnothing$. Then, there exists $\varepsilon > 0$ and $P \in \mathrm{int}(\mathrm{Sk}(X))$ such that $\mathcal{B}_\varepsilon(P) \subseteq \mathrm{int}(\mathrm{Sk}(X))$. Choose $Q \in P_C$ and $P' \in \mathcal{B}_\varepsilon(P)$ on the line segment joining P to Q making sure $P' \neq P$. Now, $|D_X(P')| \leq \mathrm{d}(P', Q)$ because Q is on the boundary of the shape. Hence, $|D_X(P')| < \mathrm{d}(P, Q)$ and $\mathcal{B}_{D_X(P')}(P') \subseteq \mathcal{B}_{\mathrm{d}(P,\, Q)}(P)$ so $\mathcal{B}_{D_X(P')}(P')$ is not maximally inscribed which contradicts Theorem A.2 (that is, Theorem 2.14). $\square$

Proposition A.4 (Proposition 2.19, p. 20). *Let X be a shape and $Q \in \partial X$. Then, there exists a $P \in \mathrm{Sk}(X)$ such that $Q \in P_C$ or Q is a sharp corner of the boundary.*

Proof. Suppose Q is not a a sharp corner. Then, recalling the definition of shape (Definition 2.6), Q is either a dull corner or ∂X is at least C^2 near Q. If there is some inscribed ball with boundary containing Q, then Theorem 2.14 guarantees that there will be a $P \in \mathrm{Sk}(X)$ such that $Q \in P_C$. Hence, it suffices to show that such a ball exists.

First, if Q is a dull corner, then the claim is trivial. On the other hand, if Q is a smooth point of the boundary, then it is a standard result that such a ball must

A.1 A region R which intersects a branch of the skeleton S. 77
A.2 See proof of Theorem 2.14. 78
A.3 Protrusion on discrete skeleton. 82

B.1 The two functions: $\sup_{t \in I(s)} \cos(t)$ and $\cos(0+s)$ are shown. 85
B.2 Configuration of ∇D_X near the skeleton. 86
B.3 Configuration maximizing δ. See proof of Lemma B.3. 87
B.4 Approximating $f''(s)$ 89

exist: any circle of radius less than $1/\kappa_B$ (where κ_B is the curvature of the boundary at Q) with center along the normal at Q and tangent to Q will do; see a standard text on Differential Geometry such as Guggenheimer [21] or DoCarmo [12]. □

Theorem A.5 (Theorem 2.24, p. 23). *Let P be a regular skeletal point and denote by Q_1 and Q_2 the two distinct points in P_C. Let $S(t)$ be the skeletal curve that that passes through P, such that $S(t_0) = P$. Then,*

(a) the angle $\angle Q_1PQ_2$ is bisected by $S'(t_0)$, i.e. the tangent to $S(t)$ at t_0; and

(b) $|r'(t_0)| = \cos\alpha$, where $r(t_0)$ is the radius function at P (i.e. $P = S(t_0)$ and $(S(t_0), r(t_0)) \in \mathbf{ST}(X)$) and $\alpha \leq \frac{\pi}{2}$ is half of $\angle Q_1PQ_2$.

Proof. [1] Let Q be one of Q_1 and Q_2. Then,

$$\mathrm{d}(Q,\ S(t_0)) = r(t_0) \quad \text{and} \quad \mathrm{d}(Q,\ S(t_0+\tau)) \geq r(t_0+\tau) \tag{A.1}$$

for $\tau \in (-\varepsilon, \varepsilon)$ and sufficiently small ε. The inequality holds because Q need not be the closest point on ∂X to $S(t_0+\tau)$. So, define $f : (-\varepsilon, \varepsilon) \longrightarrow \mathbb{R}$ as

$$\begin{aligned} f(\tau) &= \mathrm{d}(Q,\ S(t_0+\tau))^2 - r(t_0+\tau)^2 \\ &= \langle\, Q - S(t_0+\tau),\ Q - S(t_0+\tau)\,\rangle - r(t_0+\tau)^2 \end{aligned}$$

Theorem 2.21 guarantees that $S(t)$ and $r(t)$ are differentiable in some neighborhood of t_0, so decrease ε so that $f(\tau)$ becomes C^1. The derivative of $f(\tau)$ is obtained by

$$\begin{aligned} \frac{\mathrm{d}f}{\mathrm{d}\tau}(\tau) &= \left\langle\, \tfrac{\mathrm{d}}{\mathrm{d}\tau}\left(Q - S(t_0+\tau)\right),\ Q - S(t_0+\tau)\,\right\rangle - 2r'(t_0+\tau)r(t_0+\tau) \\ &= 2\,\langle\, -S'(t_0+\tau),\ Q - S(t_0+\tau)\,\rangle - 2r'(t_0+\tau)r(t_0+\tau) \end{aligned}$$

Now, the conditions in Equation A.1 translate into $f(0) = 0$ and $f(\tau) \geq 0$. Thus, $f(\tau)$ has a local extremum at zero or is constant at zero, i.e. $f'(0) = 0$. Hence,

$$\langle\, S'(t_0),\ Q - S(t_0)\,\rangle = -r'(t_0)r(t_0)$$

and since $|\langle\, S'(t_0),\ Q - S(t_0)\,\rangle| = \|S'(t_0)\|\,\|Q - S(t_0)\|\cos\alpha$, assuming $S(t)$ is parameterized by arc-length,

$$|r'(t_0)| = \cos\alpha$$

[1] This proof was adapted from [13, Theorem 6.3].

which shows the second claim. To prove the first claim, recall that the derivation holds for both Q_1 and Q_2. Since $r(t_0) = \mathrm{d}(P, Q_1) = \mathrm{d}(P, Q_2)$, 2α must be the angle $\angle Q_1 P Q_2$. □

Theorem A.6 (Theorem 2.28, p. 26). *Let $S(t)$ be a segment of a skeletal curve such that $r(t)$ is monotonically decreasing and $\|S'(t)\| = 1$. Assume further that $S(t)$ is at least C^2 and let $C(t)$ be a contour reconstruction according to Equation 2.5. Then, the curvature of the boundary segment $C(t)$, denoted $\kappa_C(t)$, is*

$$\begin{aligned} |\kappa_C(t)| &= \frac{|\alpha'(t) - \kappa_S(t)|}{\|C'(t)\|} \\ &= \left| \frac{\alpha'(t) - \kappa_S(t)}{r(t)(\alpha'(t) + \kappa_S(t)) - \sin\alpha(t)} \right| \end{aligned}$$

where $\kappa_S(t)$ is the curvature of $S(t)$.

Proof. This result is a direct consequence of the following to lemmas

Lemma A.7. *Under the assumptions of Theorem A.6,*

$$|\kappa_C(t)| = \frac{|\alpha'(t) - \kappa_S(t)|}{\|C'(t)\|}$$

where $\kappa_S(t)$ is the curvature of $S(t)$.

Proof. Denote by $\theta_S(t)$ and $\theta_C(t)$ the orientations functions of the tangents to $S(t)$ and $C(t)$ respectively. Hence, $\alpha(t + \mathrm{d}t) - \alpha(t) = \mathrm{d}\theta_S(t) + \mathrm{d}\theta_C(t)$ which implies that

$$\alpha'(t) = \kappa_S(t) + \frac{\mathrm{d}\theta_C}{\mathrm{d}t}.$$

Now, $\frac{\mathrm{d}s}{\mathrm{d}t} = \|C'(t)\|$ where $s(t)$ is the arc-length parameterization of $C(t)$. Let $g(s(t)) = t$, i.e. the inverse function of $s(t)$. The assumtions make $g(t)$ into a well-defined, well-behaved (e.g. differentiable) function. Hence,

$$\frac{\mathrm{d}g}{\mathrm{d}t} = \frac{\mathrm{d}g}{\mathrm{d}s}\frac{\mathrm{d}s}{\mathrm{d}t} = 1$$

which shows that $\frac{\mathrm{d}t}{\mathrm{d}s} = \frac{1}{\|C'(t)\|}$. Finally,

$$|\kappa_C(t)| = \left| \frac{\mathrm{d}\theta_C}{\mathrm{d}s}\frac{\mathrm{d}s}{\mathrm{d}t} \right| = \frac{|\alpha'(t) - \kappa_S(t)|}{\|C'(t)\|}$$

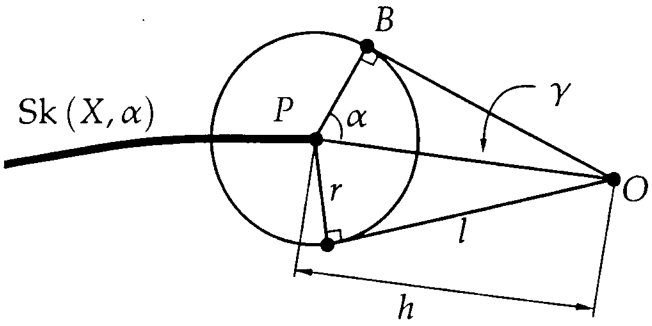

FIGURE A.3. Protrusion on discrete skeleton. The thick curve is the detected skeleton; the segment PQ is worst possible (creating most area) extension of the skeleton – the added area, then, is the portion of the triangles not in the disc.

□

Lemma A.8. *The local boundary to axis ratio, $\frac{ds}{dt}$, under the assumptions of Theorem A.6 makes sense and is given by*

$$\left(\frac{ds}{dt}\right)^2 = \|C'(t)\|^2 = \left(r(t)(\alpha'(t) + \kappa_S(t)) - \sin\alpha(t)\right)^2$$

Proof. A direct calculation from Equation 2.5 shows this. Alternatively, the result is derived by Blum and Nagel [7] in a completely different manner. □

This finishes the proof of Theorem A.6. □

Theorem A.9 (Theorem 2.31). *Let $A(X)$ denote the area of shape X and $A(X,\alpha)$ the area of the shape represented by the α-skeleton transform of X, then*

$$\frac{A(X,\alpha)}{A(X)} \leq \frac{\pi}{\pi + \tan(\alpha) - \alpha} \tag{A.2}$$

Proof. Let P be an end-point of the α-skeleton and pick any path $\gamma \in \Gamma_P$. By definition, the object angle on any point in γ is bounded above by α. Suppose that γ is a straight line as depicted on Figure A.3. Then, the triangle PBO and its reflection must contain the boundary. Therefore, the largest area of the shape not covered by the circle at P is achieved if the object angle along γ is equal to α everywhere and, consequently, the boundary is given by the two triangles as depicted. Now, the area of the triangle PBO can be obtained by noticing that $PB = r$ at P and that

$BO = r\tan(\alpha)$, which implies that $A(PBO) = r^2\tan(\alpha)/2$. Hence, the area of the boundary not covered by the disc, denoted A_{out}, is no more than $r^2\tan(\alpha) - \alpha r^2$. So,

$$\frac{A(disc)}{A_{out} + A(disc)} \leq \frac{\pi r^2}{\pi r^2 + r^2\tan(\alpha) - \alpha r^2} = \frac{\pi}{\pi + \tan(\alpha) - \alpha}$$

which shows the claim if γ is a line segment. If, on the other hand, γ_0 were not straight (e.g. with nontrivial curvature), then the area represented by γ_0 would be strictly smaller than the area represented by a straight γ. Hence, the claim also holds for such γ_0.

Finally, since the result holds for *any* $\gamma \in \Gamma_P$, it holds for the whole shape. $\square$

Appendix B

Proofs for Chapter 4

Proposition B.1 (Proposition 4.1, p. 47). *Given a shape X with no isolated points in ∂X, then there must be $\sigma_0 > 0$ such that if $\sigma < \sigma_0$ then $\mathcal{S}_\sigma\left(L_\sigma(X)\right)$ is homotopy equivalent to X.*

Proof. We want to show that every hole in X will have a unique represetative in the discrete lattice for a small enough σ. The key observation is that there are only finitely many holes in X, none of which may consist of a single point (see Definition 2.6). Thus, there is a ball of positive radius, denoted r, that fits in all holes. Consequently, if $\sigma < r/3$, then each hole will contain at least on lattice point and the claim follows. □

Lemma B.2 (Lemma 4.2, p.48; Lemma 4.7, p.50). *Let $\theta_0, \theta_1 \in \mathbb{R}$ such that $0 \leq \theta_1 - \theta_0 \leq 2\pi$. Let $E(s) : [\theta_0, \theta_1] \to [-\delta/2, \delta/2]$ for $0 \leq \delta \leq 2\pi$. Then*

$$\left| \int_{\theta_0}^{\theta_1} \cos(\alpha + E(s) + s)\ \mathrm{d}s \right| \leq \left| \int_{\theta_0}^{\theta_1} \cos(\alpha + s)\ \mathrm{d}s \right| + 2\delta$$

and

$$\left| \int_{0}^{2\pi} \cos(\alpha + E(s) + s)\ \mathrm{d}s \right| \leq \delta + 2\sin\left(\frac{\delta}{2}\right)$$

Further, if $I^{(n)}$ is the n-point approximation of the flux (see Equation 4.1), then

$$\left| I^{(n)} \right| \leq \delta + 2\sin\left(\frac{\delta}{2}\right) + \frac{4}{3}\frac{\pi^3}{n^2}\,.$$

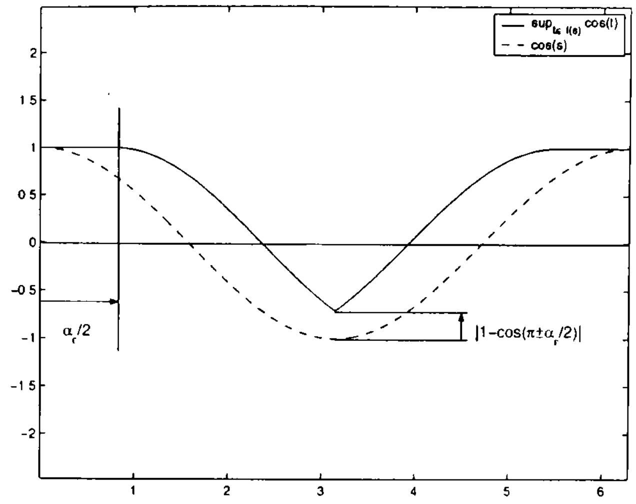

FIGURE B.1. The two functions: $\sup_{t\in I(s)} \cos(t)$ and $\cos(0+s)$ are shown.

Proof. First, if $\theta_1 - \theta_0 \le \delta$ then the claim follows since

$$\left| \int_{\theta_0}^{\theta_1} \cos(\alpha + E(s) + s)\, \mathrm{d}s \right| \le \theta_1 - \theta_0 \le \delta.$$

So, assume $\theta_1 - \theta_0 > \delta$ and let $I(s) = [\alpha + s - \delta/2, \alpha + s + \delta/2]$. We have

$$\begin{aligned} \cos(\alpha + E(s) + s) &\le \sup_{t\in I(s)} \cos(t) \\ \text{so} \quad \int_{\theta_0}^{\theta_1} \cos(\alpha + E(s) + s)\, \mathrm{d}s &\le \int_{\theta_0}^{\theta_1} \sup_{t\in I(s)} \cos(t)\, \mathrm{d}s \end{aligned}$$

Similarly, $\int_{\theta_0}^{\theta_1} \inf_{t\in I(s)} \cos(t)\, \mathrm{d}s \le \int_{\theta_0}^{\theta_1} \cos(\alpha + E(s) + s)\, \mathrm{d}s$ and therefore

$$\left| \int_{\theta_0}^{\theta_1} \cos(\alpha + E(s) + s)\, \mathrm{d}s \right| \le \left| \int_{\theta_0}^{\theta_1} \sup_{t\in I(s)} \cos(t)\, \mathrm{d}s \right|$$

or

$$\left| \int_{\theta_0}^{\theta_1} \cos(\alpha + E(s) + s)\, \mathrm{d}s \right| \le \left| \int_{\theta_0}^{\theta_1} \inf_{t\in I(s)} \cos(t)\, \mathrm{d}s \right|$$

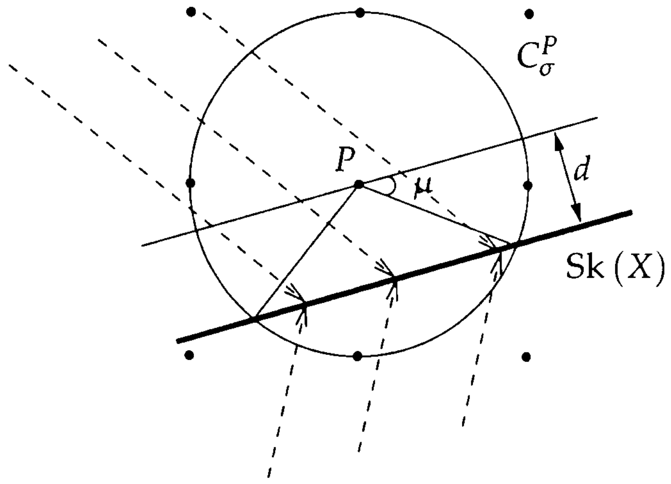

FIGURE B.2. Configuration of ∇D_X near the skeleton.

Now, for any $s \in [0, 2\pi]$ one can verify that

$$\sup_{t \in I(s)} \cos(t) = \begin{cases} 1 & \text{if } s \in [0, \delta/2] \\ \cos(\alpha + s - \delta/2) & \text{if } s \in [\delta/2, \pi] \\ \cos(\alpha + s + \delta/2) & \text{if } s \in [\pi, 2\pi - \delta/2] \\ 1 & \text{if } s \in [2\pi - \delta/2, 2\pi] \end{cases}$$

See Figure B.1 for a picture. Notice that as far as the areas under the curves are concerned, the modified cosine (sup cos) has two portions which are exactly the same as for $\cos(\alpha + s)$. The portion lacking is only that with area $\int_{-\delta}^{\delta} \cos(s) < \delta$ and it appears with negative contribution in the original cosine. The only other difference is the area of size δ around the peaks. Observing that the case for inf cos is exactly similar finishes thethe proof.

Lemma 4.2 now follows by taking $\theta_0 = 0$ and $\theta_1 = 2\pi$. Lemma 4.7 is also a consequence of this result, it provides the error bound. The $2\sqrt{3}\sin(\alpha)$ is a consequence of assuming the point to be at most $\sigma/2$ from the skeleton (see Figure B.2); thus, Equation 3.6 must account for this by adjusting the bounds on the integrals as follows:

$$\begin{aligned} \mathcal{F}_\varepsilon(P) &= \int_0^{2\pi\varepsilon} \left\langle \nabla D(C_\varepsilon^P(s)), \mathcal{N}(s) \right\rangle \, \mathrm{d}s \\ &= -\varepsilon \int_{-\mu}^{\pi+\mu} \cos(\alpha - s) \, \mathrm{d}s - \varepsilon \int_{\pi+\mu}^{2\pi-\mu} \cos(-\alpha - s) \, \mathrm{d}s \\ &= -4\varepsilon \sin(\alpha) \cos(\mu) \end{aligned}$$

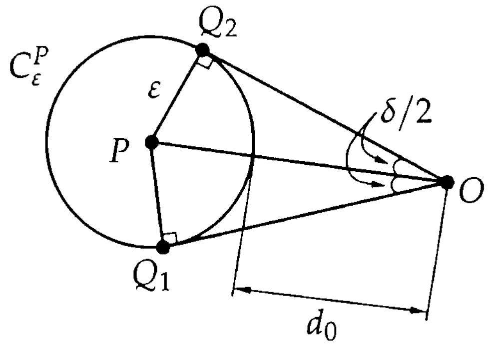

FIGURE B.3. Configuration maximizing δ. See proof of Lemma B.3.

where $\mu = \arcsin\left(\frac{d}{\varepsilon}\right)$ and $\frac{d}{\varepsilon} \leq \frac{1}{2}$, which shows the claim. □

Lemma B.3 (Lemma 4.3, p. 49). *Let X be a shape and ∂X its boundary. Let C^P_σ be a circle centered at P with radius σ. Assume that C^P_σ does not contain in its interior any skeletal points or points on ∂X. Let δ be the maximum difference in orientation of ∇D_X through C^P_σ. Then,*

$$\delta \leq 2 \arcsin\left(\frac{\sigma}{d+\sigma}\right)$$

where $d = \min\left\{\mathrm{d}\left(C^P_\sigma, \partial X\right), \mathrm{d}\left(C^P_\sigma, \mathrm{Sk}\left(X\right)\right)\right\} > 0.$

Proof. Pick any point R on C^P_σ and arrange coordinates so that the orientation of $\nabla D(R)$ is zero and so that every other ∇D vector along the circle has positive orientation. Now if $Q \in C^P_\sigma$ and $\nabla D(Q)$ does not coincide with the tangent to C^P_σ at Q, then there exists $Q' \in C^P_\sigma$, $Q' \neq Q$ such that if $B \in \partial X$ is the closest boundary point to Q, then it is also to Q'. This is because otherwise C^P_σ will have a skeletal point in its interior. Therefore, the ∇D vector with largest valued orientation δ must be such that $\delta \leq \pi$.

We now find the maximum value for δ given σ and d. Denote by Q_1 and Q_2 any pair of points on the circle such that $\langle\, \nabla D\left(Q_1\right), \nabla D\left(Q_2\right) \rangle = \cos(\delta)$. Since $\delta > 0$ (e.g. Figure B.3), suppose the lines defined by $\ell_i(t) = Q_i + t\nabla D\left(Q_i\right)$ (for $i = 1, 2$) intersect at point O. Hence, $d \leq d_0 = d(C^P_\sigma, O)$ because otherwise $O \in \partial X$ or $O \in \mathrm{Sk}\left(X\right)$ and the definition of d would not hold. It is easy to see that δ is maximized when the $\ell_i(t)$ are tangent to C^P_σ so we leave the proof to the reader. This configuration is illustrated in Figure B.3. Notice that since $PO = \sigma + d_0$, we have

$$\sin\left(\frac{\delta}{2}\right) = \frac{\sigma}{\sigma + d_0}$$

which is maximized when d_0 is minimized, i.e. when $d_0 = d$. This finishes the proof. $\square$

Lemma B.4 (Lemma 4.5, p. 49). *Let X be a shape and denote by $\kappa(t)$ the curvature function of ∂X wherever it is defined. If the boundary is made of line segments then*

$$\mathrm{NumErr}(n) \leq \frac{4}{3}\frac{\pi^3}{n^2}$$

Suppose $X \subseteq \overline{\mathcal{B}_R(P)}$, $|\kappa(t)| \leq K_1 \in \mathbb{R}$ and $|\kappa'(t)| \leq K_2 \in \mathbb{R}$. If $\frac{1}{K_1} - \frac{R}{2} > \sigma$ and $n = 2^k$ for $k \geq 2$ then

$$\mathrm{NumErr}(n) \leq \frac{4}{3}\frac{\pi^3}{n^2}\mathcal{E}\left(\frac{1}{K_1} - \frac{R}{2}\right)$$

where

$$\begin{aligned} \mathcal{E}(d) \quad &\leq \quad K_2\frac{R}{d}\,\frac{8 + 38d^2 + 25d^3 + 8d^4 + 28d + d^5}{d^5} \\ &\quad + \quad \frac{8 + 16d + 2d^3 + 10d^2}{d^5} \end{aligned}$$

Proof. We have to show that the numerical error ensuing from the n-point approximation of the AOF using the Trapezoidal Rule is as claimed. To do this, notice that the error bound for such numerical integration is given by

$$\int_{l\frac{2\pi}{n}}^{(l+1)\frac{2\pi}{n}} f(s)\mathrm{d}s = \frac{1}{12}\left(\frac{2\pi}{n}\right)^3 \sup_{s\in[l\frac{2\pi}{n},(l+1)\frac{2\pi}{n}]} |f''(s)|$$

(e.g. see Allen and Isaacson [1, p. 317]). In our case,

$$f(s) = \cos(\alpha + E(s) + s)$$

so

$$f''(s) = -\cos(\alpha + E(s) + s)\left(\frac{\mathrm{d}E}{\mathrm{d}s}(s) + 1\right)^2 - \sin(\alpha + E(s) + s)\frac{\mathrm{d}^2E}{\mathrm{d}s^2}(s).$$

Now, since $f(s)$ expresses the inner product of the normals to a circle and the corresponding value of the gradient field ∇D_X, the function $\alpha + E(s)$ actually expresses the orientation of the gradient vector at a particular point on the circle. Thus, $\frac{\mathrm{d}E}{\mathrm{d}s}$ is related to the curvature of the boundary point Q which gives rise to the gradi-

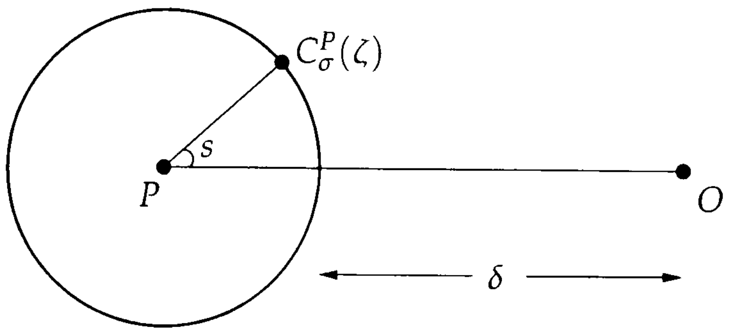

FIGURE B.4. The value of $f'(s_0)$ may be approximated by this case. O is the center of curvature for the boundary point $Q(s_0)$ corresponding to $C_\sigma^P(s_0)$ so, locally, the gradient lines converge in O. Thus, the gradient vectors at $C_\sigma^P(\zeta_0)$ are given by normalizing $O - C_\sigma^P(\zeta_0)$. $C_\sigma^P(\zeta)$ is a reparameterization for the above and such that $C_\sigma^P(\zeta_0) = C_\sigma^P(s_0)$. See proof of Lemma B.4.

ent vector with orientation $\alpha + E(s)$. If $C(t)$ is arc-length parameterization of the boundary near Q and s is as in the arc-length parameterization of $C_\sigma^P(s)$, then

$$\frac{\mathrm{d}E}{\mathrm{d}s} = \frac{\mathrm{d}C}{\mathrm{d}t}\frac{\mathrm{d}t}{\mathrm{d}s} = \kappa(t)\frac{\mathrm{d}t}{\mathrm{d}s}$$

$$\frac{\mathrm{d}^2E}{\mathrm{d}s^2} = \kappa'(t)\left(\frac{\mathrm{d}t}{\mathrm{d}s}\right)^2 + \kappa(t)\frac{\mathrm{d}^2t}{\mathrm{d}s^2}$$

Hence, if the boundary is locallly a line segment, then

$$|f''(s)| = \cos(\alpha + E(s) + s) \leq 1,$$

which shows the first part of the claim.

Now suppose that the boundary is not locally a line segment and let $d = \frac{1}{K_1} - \frac{R}{2}$. Thus, if $Q(s) \in \partial X$ is the closest point to $C_\sigma^P(s) \in \mathrm{int}\,(X)$ (the AOF is computed for the circle C_σ^P), then d is the shortest distance between $C_\sigma^P(s)$ and the center of curvature corresponding to $Q(s)$. A tight upper bound for $\frac{\mathrm{d}t}{\mathrm{d}s}$ is $\left|\frac{\mathrm{d}t}{\mathrm{d}s}\right| \leq \frac{R}{d}$.

Now, to estimate $f''(s)$, notice that all possible values for $f'(s)$ can be computed exactly given the distance δ between the circle and the center of curvature for the boundary point corresponding to $C_\sigma^P(s)$. This is because $f'(s)$ only depends that curvature and $\frac{\mathrm{d}t}{\mathrm{d}s}$, and the circle of curvature approximates it behavior locally. So,