Voronoi Skeletons: Theory and Applications

R. Ogniewicz and M. Ilg

Communication Technology Laboratory
Swiss Federal Institute of Technology ETH
CH-8092 Zurich, Switzerland

Abstract

The paper presents a novel method of robust skeletonization based on the Voronoi diagram (VD) of boundary points, which is characterized by correct Euclidean metrics and inherent preservation of connectivity. The regularization of the Voronoi medial axis (VMA) in the sense of Blum's prairie fire analogy is done by attributing each component of the VMA with a measure of prominence and stability. The resulting Voronoi skeletons (VSK) appear largely invariant with respect to typical noise conditions in the image and geometric transformations. Hierarchical clustering of the skeleton branches, the so-called skeleton pyramid, leads to further simplification of the skeleton. Several applications demonstrate the suitability of the Voronoi skeleton to higher order tasks such as object recognition.

1 Introduction

During the last decades, skeletonization or thinning has been a constant research topic. The concept of skeletonization denotes a process, which transforms a 2D object into a 1D line representation, similar to a stick figure. A concise definition of the skeleton or *medial axis* (MA) in the continuum was given by Blum [1], who postulated the well-known prairie fire analogy. In spite of its apparent simplicity, the implementation of Blum's definition in the discrete world without losing important properties such as connectivity or Euclidean metrics has turned out to be surprisingly tedious (for a survey, see [2]). Recently, an implementation of the fire front propagation in the discrete plane was proposed in [3]. However, this algorithm is based on regular metrics and therefore inherits several of the typical drawbacks of common thinning methods. Basically, we can distinguish between two substantially different variants of generic (i.e., which do not presume a very specific object representation, e.g., polygonal shapes) skeletonization algorithms: *Topological thinning* and *medial axis extraction from a distance map.*

A large class of thinning algorithms examine the topological relevance of object pixels rather than the metric properties of the shape. Typically, object pixels are repetitively tested and subsequently deleted, when-

ever their removal does not alter the topology of the thinned shape. While working fast and relatively reliably for elongated silhouettes, thinning of shapes characterized by a large diameter can often strike the human observer by counterintuitive results. The topology of the rectangular grid implies that this sort of thinning leads to regular metrics. On one hand, topological thinning can guarantee connected skeletons, on the other hand, we pay for retaining connectivity with the loss of Euclidean metrics.

The alternative approach requires the evaluation of a *distance map*. The use of regular metrics leads to very simple algorithms, but the resulting skeletons are usually incongruous with the fire front paradigm. On the other hand, efficient methods to obtain correct Euclidean distance maps have been published (e.g., [4]).

However, the next step, namely the extraction of the MA imposes severe difficulties. If we compute the skeleton as the set of endpoints of shortest intrusion paths [5], the skeleton will be characterized by a large number of redundant elements. Conversely, the computation of the centers of largest inscribed disks leads to a skeleton consisting of more or less sparsely distributed points. Therefore, other methods apply differential geometry to extract the skeleton. In that case, pseudo-Euclidean or Euclidean metrics may be used, but the problem of connectivity remains unsolved. A promising approach seemed to be the combination of thinning and a method, which explicitly extracts skeletal pixels [6]. Although it was possible to let the thinning process make the connections between disjoint actual skeleton points, the previously mentioned deficiencies of thinning algorithms influenced the quality of the results. Still, the outcome was only a rough approximation of Blum's concept. So other approaches try to complete the skeleton by tracking the ridges in the distance map (e.g., [7]). However, even in case of regular metrics, the algorithms have to handle specific situations such as the occurrence of saddle-points.

In [8], it is pointed out that a major weakness of the "symmetric (or medial) axis transform (SAT or MAT)" is its sensitivity to details of the boundary. Several methods have been proposed to avoid this deficiency. A pruning approach used in [9] is based on the analysis

of disk radii along a skeleton branch ("propagation velocity"). However, the pruned skeletons do not preserve their initial connectivity. Insensitivity to artefacts can be obtained by smoothing the boundary with a lowpass filter [10] or by blurring the shape [8]. The drawbacks of these methods are twofold: First, skeletonization has to be computed afresh for numerous levels of resolution, in order to obtain a hierarchical description. Second, postprocessing is needed to establish correspondences between segments of the MA pertaining to different levels. Still another way is to attribute each point of the MA with a measure of "prominence" of the associated boundary [11]. Consecutively, after thresholding, less essential sections of the MA are disconnected from the "backbone". However, the straightforward application of the above method to complex shapes can lead to several disjoint "backbone" sections, which intuitively should be connected. In [12] the local curvature of the boundaries of the original object and successive erosions thereof controls the preservation of potentially stable skeleton branches. Similarly, in [13], curvature extrema lead to the set of initial control points for the simulation of the "dynamic grassfire" using the *snake model*. Two questions arise with respect to the general scope of this approach. First, a ragged boundary may prematurely initiate skeleton branches, which at a later stage turn out to be irrelevant to the overall shape description. Second, the computation of the correct curvature cannot be transferred in a straightforward way to the discrete plane.

2   The Voronoi Medial Axis

Since the distance map reflects the proximity relations among boundary points, the evaluation of a distance map is closely related to the notion of proximal polygons within a point set. According to [14], such problems can be seamlessly solved with the help of the Voronoi diagram (VD) of the boundary points. The computation of a distance map can then be defined as a mapping of the Voronoi polygons onto the rectangular grid.

It is well-known that the medial axis of polygonal shapes can be obtained by computing the VD of the boundary line segments (e.g., [9], [15]). Such *Voronoi medial axes (VMA)* consist of segments of straight lines and parabola. However, any convex locus on the boundary such as the vertex of a polygon induces an additional skeleton branch. Thus the proper polygonal approximation of a shape becomes crucial for the complexity and topology of the skeleton. Unfortunately, *robust* polygonal approximation of arbitrary shapes turned out to be fairly difficult. Natural shapes require a rather large number of vertices for an accurate polygonal approximation. Numerous additional skeleton branches are thus introduced, which do not contribute essentially to the

overall representation.

Despite these obstacles, the VD of the boundary points lays the foundation of the Voronoi skeleton. In order to support symmetrical treatment of foreground and background, the course of the boundary is expressed using a symbolic description, namely as a chain of *pixel raster cracks*, i.e., elementary vectors, which separate object pixel and non-object pixel. The VD is computed based on the raster crack endpoints.

The doubly-connected-edge-list (DCEL) has been proposed as the data structure of choice for the VD in [14]. We adapt the DCEL to related data structures as the VSK by attributing the DCEL components with supplementary information, e.g., regularization data and labeling information.

The resulting VMA represents the medial axis in the sense of Blum's proposal, if each point site initiates a concentric fire front. However, the difficulty lies in the still huge number of edges, which (a) are not relevant for the basic form of the skeleton (see, e.g., Figure 3(b)) and (b) react very sensitively to even the slightest disturbations among the point sites. Consequently, the description or identification of shape by means of the complete VMA is most likely an ill-posed problem unless a regularization method is proposed.

3   Regularization
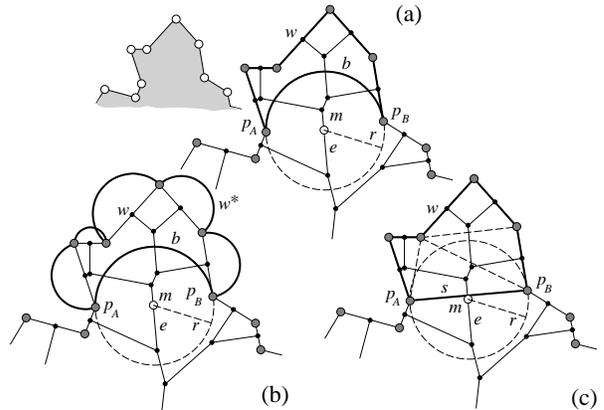
Figure 1:   Regularization of the VMA: Different Residual Functions.   (a) Potential and circularity residual.   (b) Bi-circularity residual. (c) Chord residual.

We observe that those parts of the medial axis, which lie deeply inside an object are less sensitive to changes among the boundary points than are the outer parts. Skeletal segments, which describe rather global topological relations (and therefore should be preserved) were created during the last phase of fire front propagation, when the fire fronts were approaching the quench locii from nearly opposite directions. In the case of the VMA, each Voronoi edge represents the local symmetry axis of exactly two boundary point sites, henceforth named the *anchor points* of an edge. The length of the shortest path from one anchor point to the other, measured

along the boundary B, is a strong probability indicator for the location of an edge within the shape. If this distance is large, the edge (and therefore this portion of the MA) is very likely to lie deeply inside the object. Based on these observations, the following four variants of a residual function are derived.

Potential Residual: In Figure 1(a), the anchor points of edge e are p_A and p_B. Consequently, e is attributed with the length of path w. The distance $w = \text{dist}^B(p_A, p_B)$ (superscript B denotes that the distance is measured along the boundary) for arbitrary pairs of points can be computed efficiently, if we introduce a boundary potential function $W(p)$ for every boundary point p. In order to create $W(p)$, it is necessary to track each boundary chain and assign to each point the current length of the path relative to an arbitrary origin. For example, for a closed boundary, we obtain w from

$$\text{dist}^B(p_A, p_B) = \min\{|W(p_A) - W(p_B)|, \\ L_B - |W(p_A) - W(p_B)|\}, \quad (1)$$

where L_B denotes the total length of the boundary. If p_A and p_B belong to disjoint boundary segments (typical for objects with holes), w is assigned an "infinitely" large distance value.

Since w denotes the length of that fraction of the boundary, which is spanned by the anchor points, we define a *potential residual* $\Delta R_P(e)$ (constant for every point of edge e)

$$\Delta R_P(e) \stackrel{\text{def}}{=} w = \text{dist}^B(p_A, p_B), \quad (2)$$

which attributes each Voronoi edge. Skeleton extraction then boils down to simple thresholding. Every edge with a residual value greater than a specific threshold T is assumed to be a stable part of the skeleton.

The following three variants of residual functions are closely related to the potential residual.

Circularity Residual: According to the definition of the MA, every point m of an edge e is the center of the largest inscribed disk. By comparing w with perimeter b of the disk we obtain a measure for how well the boundary segment is approximated by a circular arc (Figure 1(a)). Again, we assume that if the *circularity residual* $\Delta R_C(e, m)$ (different for every disk center m)

$$\Delta R_C(e, m) \stackrel{\text{def}}{=} w - b = \Delta R_P(e) - b \quad (3)$$

is large, the edge is likely to be a stable part of the skeleton.

Bi-circularity Residual: In order to overcome some of the shortcomings of the circularity residual such as the undue suppression of circular forms, we modify its definition by multiplying w with $\pi/2$. The so-called *bi-circularity residual* $\Delta R_B(m)$ expresses the approximation of a chain of semi-circles w^* by perimeter b of the inscribed disk (Figure 1(b)). w^* itself is an approximation of a vertex-based reconstruction of the shape,

namely by drawing the largest inscribed disk at each vertex of the VD [1]. We define $\Delta R_B(m)$ as

$$\Delta R_B(e, m) \stackrel{\text{def}}{=} \frac{2}{\pi}(\frac{\pi}{2}w - b) = \frac{2}{\pi}(\frac{\pi}{2}\Delta R_P(e) - b). \quad (4)$$

The factor $2/\pi$ assures that for elongated shapes, i.e., $b \ll w$, $\Delta R_B \approx \Delta R_C \approx \Delta R_P$.

Chord Residual: The chord residual (Figure 1(c)) is obtained by replacing perimeter b with the length s of chord $\overline{p_A p_B}$:

$$\Delta R_H(e) \stackrel{\text{def}}{=} w - s = \Delta R_P(e) - s. \quad (5)$$

Thus, $\Delta R_H(e)$ is a measure for the degree of approximation between the original shape and a subset of its *Delaunay triangulation*, the straight line dual of the VD.

It can be shown [16] that skeleton extraction by means of our residual functions preserves the connectivity of the skeletons. Thus, no additional postprocessing steps are mandatory to connect disjoint portions of the MA as in [17].

Having attributed the edges of the VMA with one of the residual functions, we need a reasonable threshold, which suppresses the effects of noisy artefacts. With the help of Figure 2, we derive such an estimate for the case of a spike-shaped disturbation along the boundary.
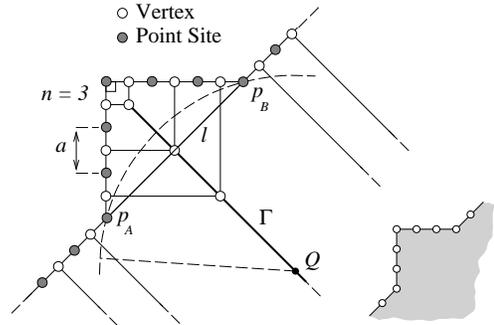
Figure 2: Estimation of a Good Threshold.

Γ denotes the spurious skeleton branch induced by the spike. We would like to suppress Γ for an arbitrarily large extent of the object. First, we obtain $l = \sqrt{2}\,na$, where n symbolizes the size of the protrusion in terms of the number of raster cracks and a denotes the length of a raster crack. Consequently, for a point Q on Γ, which lies very far away from the location of the spike, we obtain (d_e: Euclidean distance)

$$
\begin{aligned}
\lim_{d_e(Q, p_A) \to \infty} \Delta R_P(Q) &= 2na \\
\lim_{d_e(Q, p_A) \to \infty} \Delta R_C(Q) &= na(2 - \sqrt{2}) \\
\lim_{d_e(Q, p_A) \to \infty} \Delta R_B(Q) &= na(2 - \frac{2}{\pi}\sqrt{2}) \\
\lim_{d_e(Q, p_A) \to \infty} \Delta R_H(Q) &= na(2 - \sqrt{2}).
\end{aligned}
\quad (6)
$$

This set of estimates can be extended by analogous analysis of other types of artefacts.

---

[1] Note that we have introduced an additional vertex in the middle of each raster crack. This approach is useful to define a proper separation into exoskeleon and endoskeleton and to explain the meaning of ΔR_B, but not required in practical applications.
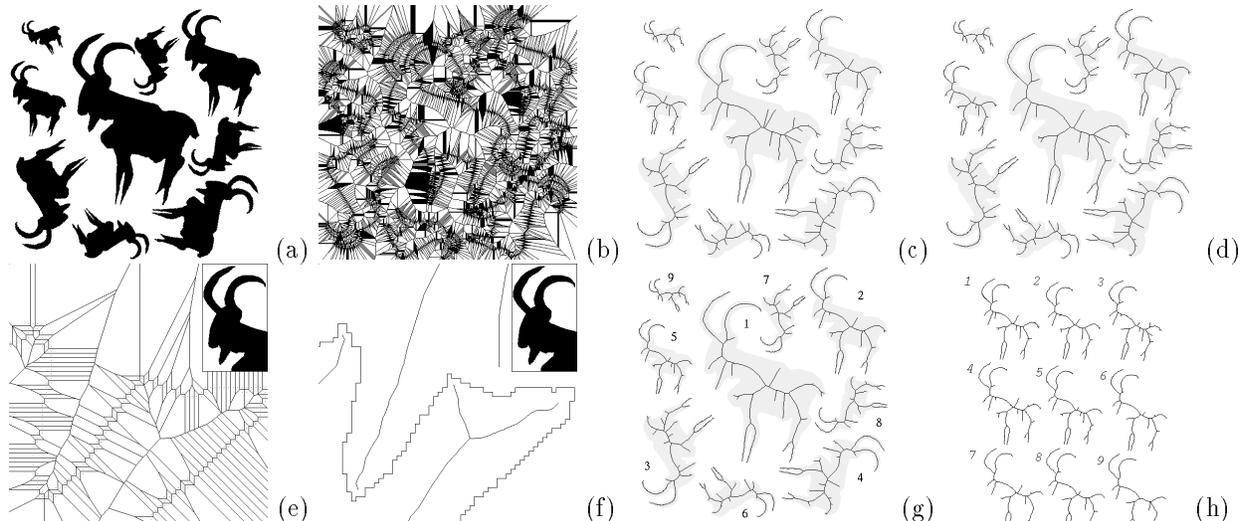
Figure 3: The Effects of Different Residual Functions. (a) Binary silhouettes (b) VD of "raster crack" endpoints. (c) ΔR_P, $T = 9.0$. (d) ΔR_H, $T = 3.0$. (e) Section of the VD. (f) Same section after regularization, ΔR_C, $T = 3.0$. (g) ΔR_C, $T = 3.0$. (h) All 9 skeletons of (g) brought to the same scale and arranged side by side.
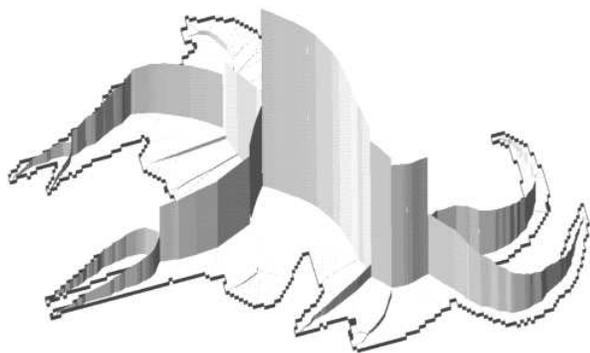
Figure 4: 3D Rendering of Circularity Residual.

In a typical configuration, an error on the order of 1 pixel results from the binarization of a nonpathological gray-valued image and a similar error is introduced by an affine transformation of the object. A useful threshold is obtained by setting the extent na of the spike to a value of $2a$ or, for the sake of security, $3a$. Since we are working with pixel raster cracks, we set $a = 1$ or $a = \sqrt{2}$ (The latter can be used to "cut off" corners). We insert these values, e.g., into Equation 6 in order to obtain a threshold estimate. The effects of thresholding are exemplified in Figures 3(c) to 3(h).

Further insight into the properties of the residual functions may be gained by using 3D visualization techniques. In Figure 4, the values of the circularity residual are rendered as the height of each Voronoi edge. The illustration is dominated by several high ridges running along the object, which can be interpreted as the *main skeleton* or the *main medial axis* of the shape. Apparently, the circularity residual discriminates very well between these portions of the skeleton, which are related to the general outline of the shape and those, which stem from tiny details.

4 Endoskeletons & Exoskeletons

The symbolic (1D) description of the boundaries results in perfect symmetry between foreground and background. Since the definition of the VMA does not refer to an explicit distinction between foreground and background, the skeletons of the foreground (endoskeletons) and of the background (exoskeletons) can be computed simultaneously. Endoskeletons describe the topology and metrics of an object, while exoskeletons depict adjacency relations between an object and its neighbors (see, e.g., Figure 6(d)). Thus the exoskeleton can be exploited for operations such as grouping of objects, edge map completion (gap closing), or path planning.

5 The Skeleton Pyramid

Albeit the pruning yields fairly stable skeletal representations, it does not solve a fundamental problem of skeletonization as will be explained next. If skeletonization should be able to handle complex compound objects, complex scenes, or objects with a significantly ragged boundary, a method to handle the still large number of skeleton branches is mandatory. A basic example of this problem is depicted in Figure 6, showing the silhouettes of three different maple leaves. After computing their Voronoi diagrams and circularity residual values, simple thresholding with $T = 3.0$ does not remove several branches, which appear to be less relevant to the general outline of the shape (Figure 6(b)). Increasing further the threshold will cause these branches to disappear eventually, although at the cost of all other branches being trimmed as well. It should be noted that neither the distance of the skeleton points from the boundary nor the length of a branch are reliable measures to decide whether a skeleton branch is
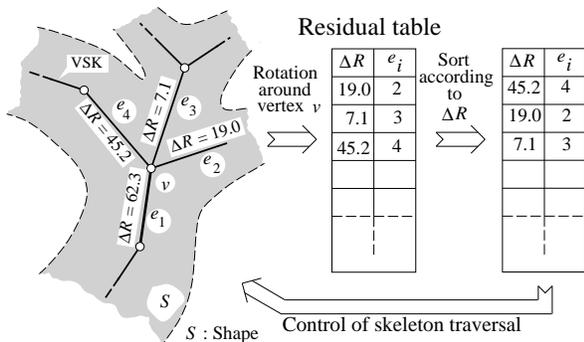
Figure 5: A Hierarchy Among the Skeleton Branches.

an important or rather spurious element.

In order to establish a topological hierarchy of the skeleton branches, the complete attributed VMA is passed on to a skeleton traversal algorithm. In the sequel, the VMA is split at the locations where it intersects with the contours into several endoskeletons and exoskeletons. Each (pruned) skeleton is attributed with a unique label, which in turn is propagated to all of its edges. Simultaneously, the largest residual value ΔR_{max} and the corresponding edge is determined. Thereafter, the skeleton is traversed a second time according to the strategy of least steep descent (Figure 5).

Let us assume that the traversal algorithm has proceeded up to edge e_1. Next, the algorithm performs a counterclockwise rotation around vertex v. So the edges e_2, e_3, and e_4 are visited in order and both their respective residual values at vertex v and their reference numbers are inserted into the so-called residual table. Afterwards, the table is sorted in descending order with respect to the residual values, so that the ordering of the edge references obeys the criterion of least steep descent. Instead of continuing the traversal according to the originally counterclockwise orientation of the edges, the process is now controlled by the sequence of edge references in the table. Referring to the values in Figure 5, the traversal will continue in edge e_4, then proceed to e_2, and finally to e_3.

Initially, the starting edge (with maximal residual value) is assigned a rank order value of 1. Let us assume that this value has been propagated up to edge e_1. The next edge being traversed, namely e_4, obtains the identical rank order value. For the subsequent edges e_2 and e_3, the behavior of the algorithm is controlled by two user defined parameters. Without these additional control mechanisms, the rank order would be incremented by one for each subsequent edge. So edges e_2 and e_3 would get assigned values of 2 and 3, respectively. The set of all edges labeled with 1 would result in a single chain without branching points. Henceforth, we denote this subset of the whole skeleton as *first order Voronoi skeleton*. Generally, for all $n > 1$ a VSK of order n denotes the union of Voronoi edges with label n and the VSK of order $n - 1$.
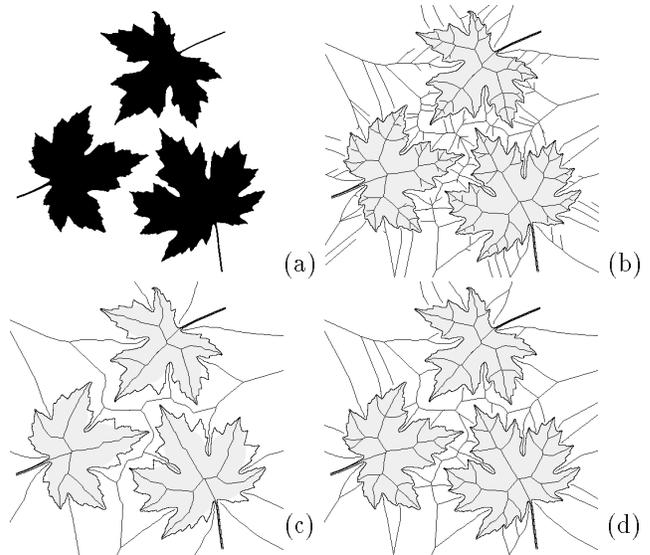
Figure 6: The Skeleton Pyramid: Two First Order VSK for Different Parameter Settings. (a) "maple leaves". (b) Voronoi skeletons, ΔR_C, $T = 3.0$. The complex boundary introduces numerous rather irrelevant branches. The shaded area (as in (c) and (d)) depicts the result of a vertex-based reconstruction of the shape, namely by drawing the largest inscribed disk at each vertex of the VSK. (c) The first order VSK of "maple leaf" shapes. Parameters $\tau_H = 0.10$, $\tau_C = 0.10$. (d) First order VSK with $\tau_H = 0.02$, $\tau_C = 0.10$.

Two new parameters are introduced to experiment with the clustering behavior of the skeleton branches, which otherwise would have been assigned different rank order values. First, parameter T_H denotes a trigger level for the build-up of the hierarchy. As long as the residual of an edge is greater than T_H, a rank order value of 1 is attached to it. It proved reasonable to adapt T_H to the individual residual maximum of each skeleton by defining $T_H = \Delta R_{max}\tau_H$. The second parameter τ_C controls the granularity of clustering within a residual table. A value of 0.1 for τ_C denotes that clusters of edges characterized by residual values within a ten percent interval of relative deviation are assigned the same rank order value. It is conceivable that better estimates for τ_H and τ_C can be derived by analyzing the distribution of the values of ΔR.

The effect of hierarchical clustering, namely the reliable segmentation of the skeleton into significant and less significant portions is depicted in Figure 6. The second order skeleton is not shown, since it is basically equivalent to the skeleton in Figure 6(b).

Notably, the outermost branches point at salient features on the boundary. Together with the complementary skeleton of the background, this observation can be employed to find a reliable simplified approximation of the shape by primitive geometric curves. An object recognition algorithm could preprocess the skeleton input by sorting all branches, which do not belong to the
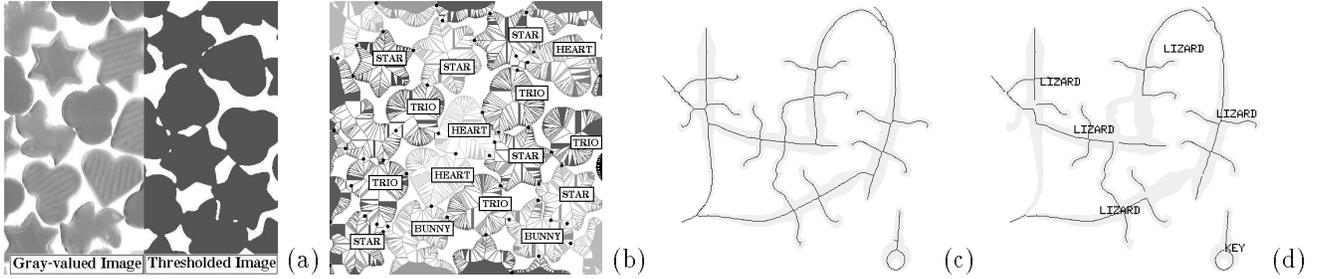
Figure 7: Object Recognition. (a) Rigid objects. Left: gray-valued image of Christmas cookies. Right: thresholded image. (b) Result of the recognition. Circularity residual. (c) Flexible, overlapping objects: ΔR_C, $T = 2.0$. (d) Result of subgraph matching. The skeletal branches shown were successfully matched with a model skeleton.

first order skeleton according to their individual maximal residual values. The identification process first analyzes the first order VSK, which was obtained by noncritical parameter settings and then proceeds by examining the branches at lower hierarchical levels.

6 Applications

6.1 Separation of Abutting or Overlapping Objects

Clusters of abutting or overlapping objects may be separated into single objects without explicit object identification. For simple geometric shapes, i.e., without characteristic narrowings, the presence of abutting or overlapping neighbors results in typical bottlenecks, which connect both neighboring shapes. Such locations may be identified as local minima of the disk radii along a skeleton branch. An application of this type of isolation was done in [18]. It helped to avoid a combinatorial explosion within an object recognition algorithm.

6.2 Object Recognition

The attributes of the VSK (disk radii, proximity information, additional contour labels) constitute a good foundation for the recognition of both geometrically rigid and flexible objects. Three implementations [18] of increasing complexity have been tested on various shapes. The fastest method only exploits the properties of the skeleton at branching points (nodes). A more sophisticated method considers the entire VMA (Figures 7(a) and (b)). The relevance and stability of each of its vertices is derived from their circularity residuals. In "critical" cases, the algorithm proceeds from the inside towards the less stable parts (boundary) and ends up in contour matching. The most elaborate version performs a subgraph matching on skeleton branches (Figures 7(c) and (d)). Each branch is attributed with additional global information like overall length, tendency of the change of disk dimensions etc.. Missing parts within the object are replaced by so–called "joker"–branches.

6.3 Extraction of Line Graphs

Image understanding is largely based on line patterns. Vectorization is therefore an important process at
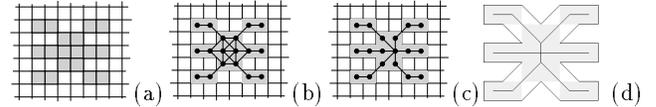
Figure 8: Graph Compilation. (a) Topologically non-deletable pixels of a line drawing. (b) Graph of (a). (c) After deletion of extra edges. (d) The VSK adequately represents the topology and metrics without additional postprocessing.

early stages. It transforms the raster pixel strings into a line graph. The limitations of the raster line generating algorithms or the inherent resolution of the underlying grid often hinder the vectorization and require costly postprocessing for structural simplification (e.g., [19]). In contrast to the above, the (vector based) VSK directly yields very appealing results (Figure 8).

6.4 Interpretation of Road Maps
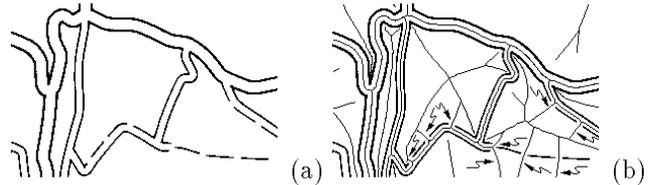
Figure 9: Understanding of Road Maps. (a) Section from a road map. (b) Corresponding first order VSK. The arrows point at skeleton branches that were used for the grouping of 'dashed' and 'para-dashed' lines.

In map recognition a scene model can not be so clear-cut but stores rather general information pertaining to all maps of interest. As a consequence, perceptual grouping is an essential activity during the transition from domain-independent image primitives to more compound entities (structured lines) representing different road types. Figure 9 shows how the information contained in both the endo- and exoskeleton of a very small portion of a map can be used effectively to assemble linearly and/or laterally structured (curved) lines (parallel lines of different width, dashed and "parallel-dashed" lines) according to the drawing rules. The interpretation of entire road maps has been described elsewhere ([20]).

Operation	Data Complex.	Computation Time
"Cracks"	11104 points	0.96 s $\approx$ 11600 points/s
VD	11104 points 31381 edges 20303 vertices	2.70 s $\approx$ 4100 points/s
ΔR_P	"	1.14 s ($\approx$ 27500 edges/s)
ΔR_C	"	2.74 s ($\approx$ 11400 edges/s)
ΔR_B	"	2.89 s ($\approx$ 10800 edges/s)
ΔR_H	"	1.66 s ($\approx$ 18900 edges/s)
Skel. Pyr. ΔR_C, $T = 3.0$		0.62 s ($\approx$ 50600 edges/s)

Table 1: Computation times on a SPARCstation-2 for the "buck" silhouettes of Figure 3 (512x512 image).

7 Summary and Outlook

The drawbacks of conventional raster-based thinning algorithms can be avoided by replacing the discrete distance map with the Voronoi diagram of boundary points. The introduction of residual functions allows to assign to each Voronoi edge a measure of prominence and stability. The resulting skeletons are characterized by true Euclidean metrics and correct topology. Further exploration of the properties of the residual functions leads to a hierarchical clustering of the skeleton components, the skeleton pyramid. Consequently, it is possible to extract a main medial axis, which represents the most significant features of the underlying shape. Further analysis of the skeleton may then turn its focus towards skeleton fragments at lower hierarchical levels. Therefore, our regularization methods and the skeleton pyramid establish a multi-resolution skeletal representation of a shape. The most important advantage of this structure is that the selection of the skeleton pertaining to a specific resolution level boils down to mere thresholding. This clearly sets it apart from other approaches such as described in [10] and [8]. Voronoi skeletons were already used in several applications as shown in Section 6.

The performance figures in Table 1 illustrate that the evaluation and intelligent handling of complex data structures such as Voronoi skeletons is not necessarily a time consuming task.

The notion of VSK has been already extended to nonclosed object boundaries of arbitrary complexity. Therefore, it is possible to pass the (symbolic rather than pixel-oriented) output of an edge detector to the VSK apparatus. Moreover, the Voronoi skeletons are eminently suitable to derive a semantic description of objects and object relations. The ongoing research is directed towards this goal and the results will be the topics of following publications.

References

[1] H. Blum, "A transformation for extracting new descriptors of shape," in *Models for the Perception of Speech and Visual Form* (W. Wathen-Dunn, ed.), Cambridge MA: MIT Press, 1967.

[2] R. W. Smith, "Computer processing of line images: A survey," *Pattern Recognition*, vol. 20, no. 1, pp. 7–15, 1987.

[3] Y. Xia, "Skeletonization via the realization of the fire front's propagation and extinction in digital binary shapes," *IEEE PAMI*, vol. 11, no. 10, pp. 1076–1086, 1989.

[4] L. Vincent, "Exact euclidean distance function by chain propagation," in *Proc. Computer Vision and Pattern Recognition, Maui, Hawaii*, pp. 521–525, 1991.

[5] U. Montanari, "A method for obtaining skeletons using a quasi-euclidean distance," *JACM*, vol. 15, no. 4, pp. 600–624, 1968.

[6] E. R. Davies and A. P. N. Plummer, "Thinning algorithms: a critique and a new methodology," *Pattern Recognition*, vol. 14, pp. 53–63, 1981.

[7] C. Arcelli and B. S. di Baja, "A width-independent fast thinning algorithm," *IEEE PAMI*, vol. 7, no. 4, pp. 463–474, 1985.

[8] S. M. Pizer, W. R. Oliver, and S. H. Bloomberg, "Hierarchical shape description via the multiresolution symmetric axis transform," *IEEE PAMI*, vol. 9, no. 4, pp. 505–511, 1987.

[9] U. Montanari, "Continuous skeletons from digitized images," *JACM*, vol. 16, no. 4, pp. 534–549, 1969.

[10] A. R. Dill, M. D. Levine, and P. B. Noble, "Multiple resolution skeletons," *IEEE PAMI*, vol. 9, no. 4, pp. 495–504, 1987.

[11] S.-B. Ho and C. Dyer, "Shape smoothing using medial axis properties," *IEEE PAMI*, vol. 8, no. 4, pp. 512–520, 1986.

[12] S. Riazanoff, B. Cervelle, and J. Chorowicz, "Parametrizable skeletonization of binary and multi-level images," *Pattern Recognition Letters*, vol. 11, pp. 25–33, 1990.

[13] F. Leymarie and M. D. Levine, "Simulating the grassfire transform using an active contour model," *IEEE PAMI*, vol. 14, no. 1, pp. 56–75, 1992.

[14] F. P. Preparata and M. I. Shamos, *Computational Geometry*. Texts and Monographs in Computer Science, New York: Springer-Verlag, second ed., 1990.

[15] D. Lee, "Medial axis transformation of a planar shape," *IEEE PAMI*, vol. 4, no. 4, pp. 363–369, 1982.

[16] R. Ogniewicz, "Regularization of the Voronoi medial axis," Tech. Rep. 127, IKT/Image Science Lab, ETH, Zürich, Switzerland, 1991.

[17] M. P. Martinez-Perez, J. Jimenez, and J. L. Navalon, "A thinning algorithm based on contours," *CVGIP*, vol. 39, pp. 186–201, 1987.

[18] U. Kienholz and R. Ogniewicz, "Isolation und Identifikation einander berührender und/oder überlappender Teile in Binärbildern," diploma thesis, ETH Zürich, Switzerland, 1987.

[19] S. Suzuki, "Graph-based vectorization method for line patterns," in *Proc. Computer Vision and Pattern Recognition*, pp. 616–621, 1988.

[20] M. Ilg, "Knowledge-based understanding of road maps and other line drawings," in *Proc. 10th ICPR 90, Atlantic City*, pp. 282–284, 1990.