

Topologically Correct Image Segmentation Using Alpha Shapes

Peer Stelldinger, Ullrich Köthe, and Hans Meine

University of Hamburg, 22527 Hamburg, Germany

Abstract. Existing theories on shape digitization impose strong constraints on feasible shapes and require error-free measurements. We use Delaunay triangulation and α -shapes to prove that topologically correct segmentations can be obtained under much more realistic conditions. Our key assumption is that sampling points represent object *boundaries* with a certain maximum error. Experiments on real and generated images demonstrate the good performance and correctness of the new method.

1 Introduction

A fundamental question of image analysis is how closely a computed image segmentation corresponds to the underlying real-world partitioning. Existing geometric sampling theorems are limited to binary partitionings, where the plane is split into (not necessarily connected) fore- and background components. In this case, the topology of the partition is preserved under various discretization schemes when the original regions are r -regular and the sampling grid has a maximum pixel radius of $r' \leq r$ [7, 9]. By making slightly stronger assumptions, this property is preserved when the shapes are blurred by a disc or square of radius p prior to discretization [6, 10] or when regions may have corners [11].

However, these theorems have two important limitations: they are not applicable when there are more regions than just fore- and background, and they do not make any predictions about the consequences of measurement errors. One reason for these limitations is the assumption of a fixed sampling grid. We are dropping this assumption in favour of *adaptive sampling* where sampling points are placed roughly along the contour of the regions to be segmented.

Our treatment of adaptively placed sampling points is inspired by research on laser range scanning. Here, a number of isolated sampling points is scattered over the surface of the object of interest, and the task is to reconstruct the surface from the set of points. A successful solution of this problem is the concept of α -shapes [4, 5]. Under certain conditions, an α -shape is homeomorphic or at least homotopy equivalent to the desired object surface. By applying this idea to the problem of image segmentation, we are able to derive a new condition on object shape that ensures homotopy equivalence of the digital segmentation with the original analog plane partitioning. This means in particular that there is a 1-to-1 mapping between the computed and the ground-truth regions. By imposing slightly stronger requirements on region shape, these properties can even be guaranteed when the segmentation is subject to measurement errors.

2 Preliminaries

We consider the task of reconstructing a partition of the Euclidean plane from a sampled representation. The plane partition to be recovered is defined as follows:

Definition 1. A partition of the plane \mathbb{R}^2 is defined by a finite set of points $P = \{p_i \in \mathbb{R}^2\}$ and a set of pairwise disjoint arcs $A = \{a_i \subset \mathbb{R}^2\}$ such that every arc is a mapping of the open interval $(0, 1)$ into the plane, the start and end points $a_i(0)$ and $a_i(1)$ are in P (but not in a_i). The union of the points and arcs is the boundary of the partition $B = P \cup A$, and the regions $R = \{r_i\}$ are the connected components (maximal connected sets) of the complement of B .

The partition is called *binary* when we can assign two labels (foreground and background) to the regions such that every arc is in the closure of exactly one foreground and one background region. A binary partition is called *r -regular*, when at every boundary point there exist two osculating discs of radius r which are entirely in the foreground and background respectively [9, 10]. This implies that regions cannot have corners, and junctions of three or more regions are impossible. These restrictions are somewhat relaxed by the notion of *r -halfregular partitions*, where an osculating r -disc must exist at least in the foreground or the background, and the number of regions must not change under either morphological opening or closing with a disc of radius $\leq r$ [11]. Corners are now possible, but the partition is still binary and has no junctions. The two notions of r -regularity and r -halfregularity have been central to all existing geometric sampling theorems. In this paper, the class of feasible plane partitions is extended as follows:

Definition 2. A plane partition is *r -stable* when its boundary can be dilated with a closed disc of radius s without changing its homotopy type for any $s \leq r$.

In other words, we can replace an infinitely thin boundary with a strip of width $2r$ such that the number and enclosure hierarchy of the resulting regions is preserved. In particular, “waists” are forbidden, whereas junctions are allowed, see Fig. 1. This includes r -regular and r -halfregular partitions, but also allows non-binary partitions and junctions and models real images much better. In particular, *polygonal partitions* (all arcs are straight lines) are always r -stable for some sufficiently small r . Unfortunately, the traditional way of proving a geometric sampling theorem (using a fixed grid at arbitrary position and angle, in connection with subset or supercover digitization) does not work for these partitions because topological equivalence cannot be guaranteed in general. Therefore, we consider another approach to digitization: we approximate the *boundary* of the partition with a finite set of *adaptively placed* sampling points. The sampling points are selected somehow “near” the boundary. We formalize this as follows:

Definition 3. A finite set of sampling points $S = \{s_i \in \mathbb{R}^2\}$ is called a (p, q) -sampling of the boundary B when the distance of every point $b \in B$ to the nearest point in S is at most p , and the distance of every point $s \in S$ to the nearest point in B is at most q . The elements of S are called *edgels*. The sampling is said to be *strict* when all edgels are exactly on the boundary, i.e. $q = 0$.




Fig. 1. The homotopy type of an r -stable plane partition does not change when dilated with a disc of radius of at most r (light gray), while dilations with bigger radius (dark gray) may connect different arcs as marked by the circle.

The Hausdorff distance between the boundary and its sampling is $d_H(S, B) = \max(p, q)$. Obviously, $q < p$ is required because S is finite. Non-zero edgel shifts $q > 0$ can be caused by systematic or statistical measurement errors. Edgels may be determined in various ways (section 4), but this only matters in so far as it determines the accuracy of the sampling, i.e. the values of p and q . Once computed, we consider edgels as isolated points that somehow define the digital boundary and connect them by means of the Delaunay triangulation:

Definition 4. *The Delaunay triangulation D of a set of points S is the set of all triangles formed by triples $t \subset S$ such that the open circumcircle of every triangle does not contain any point of S . If the points are in general position, the Delaunay triangles, their edges and corners (also denoted as 2-, 1- and 0-cells in this context) form a uniquely defined, connected simplicial complex. The union of all cells $|D| = \bigcup_{c \in D} c$ is called the polytope of D .*

In order to approximate the boundary of the partition, we want to remove those edges and triangles from the Delaunay triangulation that are not related to the boundary. A useful subset is defined by the α -complex introduced in [4]:

Definition 5. *The α -complex D_α of a set of points S is the subcomplex of the Delaunay triangulation D of S which contains all cells c such that (a) the radius of the smallest open circumcircle of c is smaller than α , and this circle contains no point of S , or (b) an incident cell c' with higher dimension is in D_α .*

The polytope $|D_\alpha|$ is called α -shape. Since cells are removed from the Delaunay triangulation, the α -complex has holes which hopefully correspond to the regions we are trying to segment. In order to determine when this is the case, the following theorem is of fundamental importance (the proof can be found in [5]):

Theorem 1 (Edelsbrunner). *The union of closed α -discs with centers at the points $s_i \in S$ covers $|D_\alpha|$, and the two sets are homotopy equivalent.*

Consequently, the α -shape $|D_\alpha|$ is homotopy equivalent to the original plane partition if and only if the dilation of the edgels with α -discs is homotopy equivalent to the boundary of the partition. This requirement is indeed fulfilled in certain situations: In [1] it is proved that $|D_\alpha|$ is even homeomorphic to B if B is the boundary of an r -regular set with $p < \alpha < r$ and $q = 0$. Unfortunately, this no longer applies when the original partition is not r -regular and/or the edgels are not exactly on the original boundary. Fig. 2 shows an example where the r -dilation of the boundary is homotopy equivalent to the boundary (i.e. the partition is r -stable), but the dilation of the edgels is not. The rest of the paper is devoted to the question what can be said under these more general conditions.




Fig. 2. The α -dilation (a) of the boundary of an α -stable plane partition may not be homotopy equivalent to the union (b) of the α -discs centered at the edgels. Thus the α -shape (c), which is always homotopy equivalent to the union of discs (b), may contain unwanted holes consisting of Delaunay triangles of radius greater than α . Thus there exists an α -disc centered in the hole which does not cover any edgel, as shown in (c).

3 Segmentation with Alpha-Shapes


Since holes of the α -complex do not necessarily correspond to regions of the original plane partition, we must characterize these holes in more detail:

Definition 6. Consider the Delaunay triangulation D of a point set S and the complement $D_\alpha^C = \mathbb{R}^2 \setminus |D_\alpha|$ of the corresponding α -polytope with $\alpha > 0$. A connected component of D_α^C is called an α -hole of $|D_\alpha|$. When the radius of the circumcircle of the largest Delaunay triangle in an α -hole's closure is at least $\beta \geq \alpha$, we speak of an (α, β) -hole.

For simplicity, we also use the term “hole” for the component which contains the infinite region. It is an (α, β) -hole for arbitrary large β . It follows from theorem 1 that there is a 1-to-1 relation between α -holes and the holes in the union of α -discs around the edgels. A similar relationship exists for (α, β) -holes:

Lemma 1. An α -hole h is an (α, β) -hole if and only if it contains a point v whose distance from the nearest edgel is at least β .

Proof. **I** ($d_H(v \in h, S) \geq \beta \Rightarrow h$ is an (α, β) -hole): when v is in the infinite region, the claim follows immediately. Otherwise, v is contained in some Delaunay triangle. By assumption, the corners of this triangle must have distance $\geq \beta$ from v . Hence, the triangle's circumradius must be at least β and the claim follows.



II (h is an (α, β) -hole $\Rightarrow \exists v \in h$ with $d_H(v, S) \geq \beta$): by assumption, the closure of h contains a Delaunay triangle t with circumradius of at least β . Consider the center v of its circumcircle. If it is within the triangle t , it is also in h and the claim follows. Otherwise, it is at least in some (α, β) -hole, and we must prove that t is in the same hole. Suppose to the contrary that v and t are in different α -holes. Then there exists a Delaunay triangle t' or a single edge e between t and v whose smallest circumcircle is smaller than α . The corners of t' or e cannot be inside t 's circumcircle since it is a Delaunay triangle. Neither t' nor e can contain v because their circumcircle radius would then be at least β .

Now consider the illustrated triangle p_1, p_2, p_3 and its circumcircle (gray) with center v . The points p_4 and p_5 are the end points of e or of one side of t' . Their distance $|p_4p_5|$ must be greater than $|p_1p_3|$. Consequently, any circumcircle with radius $\leq \alpha$ (dashed) around p_4 and p_5 contains t , contrary to the condition (imposed by the definition of an α -complex) that it must not contain any other edge. The claim follows from the contradiction. \square

Even for optimally chosen α , the α -complex does not necessarily reconstruct the homotopy type of the original boundary, since it may contain too many holes (see Fig. 2). This can be “repaired” by identifying (α, β) -holes:

Definition 7. An (α, β) -boundary reconstruction from an edgel set S is defined as the union of the polytope $|D_\alpha|$ with all α -holes of D_α that are not (α, β) -holes.

In other words, surplus holes are simply “painted over”, and (α, β) -boundary reconstruction essentially amounts to hysteresis thresholding on the triangle size of a Delaunay triangulation. The following theorem shows that exactly the desired holes survive when α and β are properly chosen.

Theorem 2 (Boundary Sampling Theorem). Let \mathcal{P} be an r -stable plane partition, and S a (p, q) -sampling of \mathcal{P} 's boundary B . Then the (α, β) -boundary reconstruction \mathcal{R} defined by S is homotopy equivalent to B , and the (α, β) -holes of \mathcal{R} are topologically equivalent to the regions r_i of \mathcal{P} , if

1. $p < \alpha \leq r - q$
2. $\beta = \alpha + p + q$
3. every region r_i contains an open γ -disc with $\gamma \geq \beta + q > 2(p + q)$.

Proof. Let U be the union of open α -discs centered at the points of S . Furthermore, let $B^\oplus = B \oplus \mathcal{B}_{\alpha+q}^o$ be the dilation of B with an open $\alpha + q$ -disc, and $r_i^\ominus = r_i \ominus \mathcal{B}_{\alpha+q}$ the erosion of region $r_i \in \mathcal{P}$ with a closed $(\alpha + q)$ -disc.

- According to the definition of a (p, q) -sampling, the dilation of B with a closed q -disc covers S . Consequently, B^\oplus covers U . Therefore, U cannot have fewer connected components than B^\oplus . B^\oplus has as many components as B due to r -stability of \mathcal{P} . Conversely, since $\alpha > p$, every open α -disc around a point of S intersects B , and the union U of these discs covers B . It follows that U cannot have more components than B . The number of components of B and U is thus equal. Due to homotopy equivalence of U and $|D_\alpha|$ (theorem 1), this also holds for the components of $|D_\alpha|$.
- Since \mathcal{P} is r -stable with $r \geq \alpha + q$, each r_i^\ominus is a connected set with the same topology as r_i . The intersection $r_i^\ominus \cap B^\oplus$ is empty, and r_i^\ominus cannot intersect $|D_\alpha| \subset U \subset B^\oplus$. Hence, r_i^\ominus is completely contained in a single α -hole of $|D_\alpha|$.
- Due to condition 3, r_i contains a point whose distance from B is at least $\gamma = \beta + q$. Its distance from S is therefore at least $\gamma - q = \beta$. Due to lemma 1, the α -hole which contains r_i^\ominus is therefore also an (α, β) -hole.
- Since B^\oplus covers U and U covers B , no (α, β) -hole can intersect both r_i^\ominus and r_j^\ominus ($i \neq j$). It follows from this and the previous observation, that every region r_i can be mapped to exactly one (α, β) -hole which will be denoted h_i .

- An α -hole that does not intersect any region r_i^\ominus must be completely contained within B^\oplus . Every point $v \in B^\oplus$ has a distance $d < \alpha + q$ to the nearest point of B . In turn, every point in B has a distance of at most p to the nearest point in S . Hence, the distance from v to the nearest point of S is $d' < \alpha + p + q = \beta$. According to lemma 1, this means that an α -hole contained in B^\oplus cannot contain a triangle with circumradius β and cannot be an (α, β) -hole.
- The previous observation has two consequences: (i) All holes remaining in \mathcal{R} intersect a region r_i^\ominus . Therefore, the correspondence between r_i and h_i is 1-to-1, and B and $|\mathcal{R}|$ enclose the same number of regions. (ii) All differences between \mathcal{R} and D_α (i.e. all Delaunay cells re-inserted into \mathcal{R}) are confined within B^\oplus . This implies that $|\mathcal{R}|$ cannot have fewer components than B^\oplus and B . Since all re-inserted cells are incident to D_α , $|\mathcal{R}|$ cannot have more components than $|D_\alpha|$, which has as many components as B (see first observation). Hence, B and $|\mathcal{R}|$ have the same number of components.
- Consider the components of the complement $(r_i^\ominus)^C$ and recall that r_i^\ominus is a subset of both r_i and h_i for any i . Since B and $|\mathcal{R}|$ have the same number of components, it is impossible for h_i^C to contain a cell that connects two components of $(r_i^\ominus)^C$. This means that the sets r_i^C and h_i^C have the same number of components. This finally proves the topological equivalence of r_i and h_i , and implies homotopy equivalence of B and $|\mathcal{R}|$. \square

If there exists no r fulfilling all conditions of theorem 2 for a given plane partition (or if the chosen α is too big), topology preservation is no longer guaranteed. Very small regions may get lost in the reconstruction. A region that is split into two or more parts by an s -erosion (i.e. has an s -waist) with $s < \alpha$ may also be split in the reconstruction. In case of very small waists, i.e. when $s + 2p + 2q \leq \alpha$, this is even guaranteed to happen. Thus, we can still apply our sampling theorem: we modify the original plane partition by connecting the different sides of small waists by a new arc. When the new partition fulfills our requirements, the modified topology is preserved, and the difference between the modified reconstruction and the original plane partition is well defined, see the second column of Fig. 3. When a thick boundary representation is undesirable, we apply *topology-preserving thinning*. An edge in the (α, β) -boundary reconstruction is called *simple* if its removal does not change the topology of the reconstructed regions. Simple edges always bound an (α, β) -hole on one side and a triangle in the boundary reconstruction on the other. Thinning removes simple edges until none are left:

1. Find all simple edges of the given (α, β) -boundary reconstruction and put them in a priority queue (the sorting is discussed below).
2. As long as the queue is not empty: Remove the topmost edge in the queue from the boundary reconstruction when it is still simple (it may have lost this property after removal of other edges). Put the edges in the triangle of the removed edge in the queue if they have now become simple.

As far as region topology is concerned, the ordering of the edgels in the priority queue is arbitrary. For example, we can measure the contrast (image gradient) and remove weak edges first. Ordering by edge length is particularly interesting:

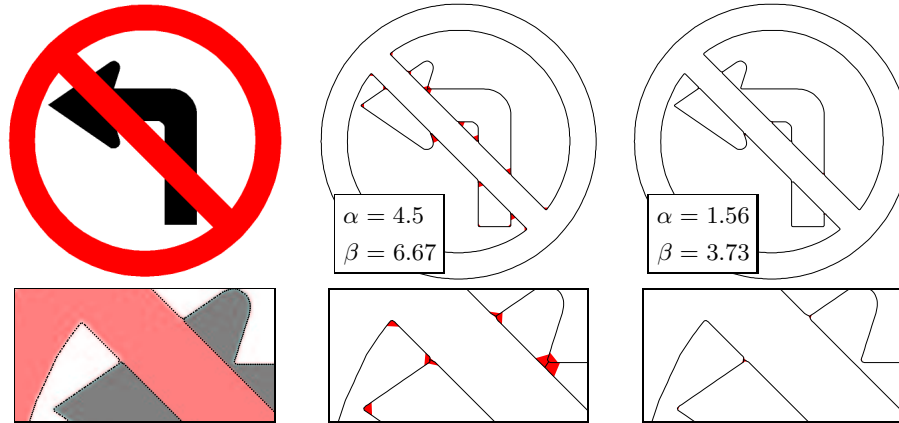


Fig. 3. Reconstructions before (red and black) and after (black only) thinning (note edges in lower left image). Connectivity errors can occur when α is too big (center).

Definition 8. A (not necessarily unique) minimal boundary reconstruction is obtained from an (α, β) -boundary reconstruction by means of topology-preserving thinning where the longest edges are removed first.

The resulting boundaries are illustrated in Fig. 3. Since region topology is preserved, a minimal boundary reconstruction is homotopy equivalent to B . The two boundaries do not in general have the same topology, because the reconstruction may contain short edges, which end in the interior of a region.

Since a minimal boundary reconstruction can be shown to be a shortest possible one with correct topology, surviving edges connect edgels closest to each other. Neighboring edgels optimally align on the thinned boundary. The length d_{\max} of the longest surviving edge is a measure of the density of the boundary sampling. The maximum distance p from the true boundary to the nearest edgel may be much larger than $d_{\max}/2$ if the displacements of neighboring edgels are highly correlated. This often occurs in practice: for example, Canny edgels along a circular arc are consistently biased toward the concave side of the curve. An (α', β) reconstruction of the edgel set with $\alpha' = d_{\max}/2 + \epsilon < p$ and arbitrarily small ϵ is still correct in the sense of theorem 2: since a minimal reconstruction is a subset of the (α', β) reconstruction, no true regions can get merged. Since $\alpha' < \alpha$, no region can get lost, and since β remains unchanged, no additional holes can be created. In fact, $\beta' = \alpha' + p + q < 2p + q$ would have been sufficient.

We found experimentally that undesirable holes (α -holes that are not (α, β) -holes) are actually quite rare, and their largest triangles are hardly ever as large as the maximal possible circumradius β allows. Therefore, an (α', β') -boundary reconstruction with β' even smaller than $\alpha' + p + q$ often produces the correct region topology. We are currently investigating the conditions which permit weaker bounds. This is important, because a smaller β leads to a correspondingly reduced γ , i.e. the required size of the original regions is reduced, and more difficult segmentation problems can be solved correctly.




Fig. 4. (a) Where the boundary intersects the dual grid, the nearest sampling points form the *grid intersection digitization*. (b) The *supercover digitization* contains all sampling points whose pixel facets intersect the arc. (c) *Canny's algorithm* produces subpixel-accurate edgels from gray scale images.

4 Application to Sampling and Segmentation Schemes

In theorem 2, p and q are assumed given. We now make their meaning and consequences more intuitive, by computing or estimating them for common sampling and segmentation schemes. Let's first look at grid intersection digitization:

Definition 9. Consider a plane partition \mathcal{P} with boundary B and a square grid. Compute all intersection points of B with the grid lines (i.e. with the lines connecting 4-adjacent sampling points) and round their coordinates to the nearest sampling point. The set of edgels thus defined is called grid intersection digitization of B , see Fig. 4a.

For simplicity, let the grid size (i.e. the smallest distance from one sampling point to another) be unity. When each component of B crosses at least one grid line, the distance p of any point of B to the nearest selected grid point is less than $\sqrt{2}$, and the distance q of any grid intersection to its rounded coordinate cannot exceed $1/2$. Inserting this into the conditions of theorem 2, we get $\alpha \geq \sqrt{2}$, $r \geq \sqrt{2} + \frac{1}{2}$, $\beta \geq 2\sqrt{2} + \frac{1}{2} \approx 3.3$, and $\gamma \geq 2\sqrt{2} + 1 \approx 3.8$. However, the worst case configurations giving rise to the values of β and γ in the theorem cannot actually occur in a square grid because Delaunay edges between grid points cannot have arbitrary length. It can be shown that the largest circumradius in an undesirable α -hole is below $\sqrt{34} \approx 2.9$, so that $\gamma \approx 3.4$ (circle area 37 pixels) is sufficient.

Generally the grid intersection digitization of a connected curve is an 8-connected digital curve. It is identical to Bresenham's digital straight line in case of a straight arc. Moreover the grid intersection digitization is a subset of the supercover digitization on a square grid, which produces a 4-connected digital curve for any connected curve:

Definition 10. Let \mathcal{P} be a plane partition with boundary B and \mathcal{G} a finite set of sampling points such that the Voronoi cells of \mathcal{G} have a radius of at most g . The supercover digitization of B is the set of all sampling points whose Voronoi cell intersects B , see Fig. 4b.

The constraint on the size of the Voronoi cells implies that $p = g$ and $q < g$. Hence, $\alpha > g$, $r > 2g$, $\beta > 3g$ and $\gamma > 4g$ are required. For example, in a unit square grid we have $q < p = \sqrt{2}/2$ and $\gamma > 2\sqrt{2} \approx 2.8$. Thus, the supercover digitization imposes weaker constraints on the original plane partition \mathcal{P} than the grid-intersection digitization. This is mainly due to the denser sampling of




Fig. 5. The *interpixel boundary* (dashed) can be extracted from the subset digitization (a). It includes both the *midcrack digitization* (b) and the *endcrack digitization* (c).

the boundary (smaller spacing of the edgels) in the former. As stated in [8], the supercover digitization is a Hausdorff discretization, i.e. a set of sampling points which minimizes the Hausdorff distance $\max(p, q)$ to the boundary B . Thus the given bounds for α , β and γ are sufficient for all Hausdorff discretizations. Another interesting question is what can be said about region based digitization:

Definition 11. Let \mathcal{P} be a plane partition with regions $R = \{r_i\}$ and \mathcal{G} a finite set of sampling points such that the Voronoi cells (i.e. the pixels) of \mathcal{G} have a radius of at most g . The subset digitization \hat{r}_i of region r_i is the union of all Voronoi cells whose sampling point is in r_i , see Fig. 5a. The union of the boundaries of all \hat{r}_i is called the *interpixel boundary*. A boundary digitization scheme where all edgels are on the interpixel boundary \mathbb{B} is an *interpixel digitization*. Two examples are the *midcrack digitization* (Fig. 5b) where the center points of all pixel edges inside \mathbb{B} are chosen as edgels, and the *endcrack digitization* (Fig. 5c) where all pixel corner points inside \mathbb{B} are used.

Thus, *boundary-based* digitizations like endcrack and midcrack digitization can be derived from the *region-based* subset digitization. While the maximal distance q of any edgel to the nearest boundary point cannot exceed g , the distance p from any boundary point to the nearest edgel can be arbitrary large: An r -stable, but non-binary plane partition is never r -regular. Consequently, \hat{r}_i is generally not topologically equivalent to the closure of r_i and may even be disconnected. The distance of the components of \hat{r}_i may approach the diameter of r_i when r_i has a long narrow spike. Obviously, this is not a useful practical bound for the value of p . We need a restriction that is stronger than r -stability, but weaker than r -regularity and which prevents these undesirable spikes:

Definition 12. Let \mathcal{P} be a plane partition with boundary B . We say two points $x_1, x_2 \in B$ delimit a (θ, d) -spike, if the distance from x_1 to x_2 is at most d and if every path on B from x_1 to x_2 contains at least one point with $\angle x_1 y x_2 < \theta$. We say that \mathcal{P} has no (θ, d) -spikes if for any pair of boundary points $x_1, x_2 \in B$ with distance of at most d , there exists a path $Y \subset B$ between x_1 and x_2 such that $\angle x_1 y x_2 \geq \theta$ for all points $y \in Y$.

Intuitively, two points delimit a (θ, d) -spike, if the shortest boundary path between them does not differ too much from a straight line, i.e. it lies inside the shaded region in Fig. 6.

Note that r -regular partitions have no (θ, d) -spikes for $d \leq r$ and $\theta = 2 \arctan(d/(2r - \sqrt{4r^2 - d^2}))$ (e.g. for $\theta = 90^\circ, 60^\circ$ we get $d = r$ and $d = \sqrt{3}r$




Fig. 6. Any point which encloses an angle of at most θ with x_1 and x_2 must lie inside the shaded region. The shown y is the one with the maximal distance to the nearer one of x_1 and x_2 . Thus there is a path from x_1 to x_2 inside the shaded region and each of its points has a distance of at most $d/(2 \sin \frac{\theta}{2})$.

respectively). By sampling dense enough one can enforce the angles to be arbitrarily flat. But in general, absence of (θ, d) -spikes does not imply r -stability, so we will require both. The fact that the boundary cannot be too far away from the edgels can be used for estimating p and q , e.g. in the midcrack and endcrack digitization case:

Theorem 3. *Let G be a square grid with sample distance h (pixel radius $g = h/\sqrt{2}$), and let \mathcal{P} be a plane partition such that every region $r_i \in \mathcal{P}$ contains a closed g -disc and the boundary B has no (θ, d) -spikes. Then the endcrack digitization is a (p, q) -boundary sampling with $q = h/\sqrt{2}$ and $p = q + (\frac{h}{2} + q) / \sin \frac{\theta}{2}$, provided that $h \leq d/(1 + \sqrt{2})$, and the midcrack digitization is a (p, q) -boundary sampling with $q = \frac{h}{2}$ and $p = q + (\frac{h}{2} + q) / \sin \frac{\theta}{2}$, provided that $h \leq \frac{d}{2}$.*

Proof. First, we prove the bounds on q . Let x, y be two 4-adjacent square grid points. Their common pixel edge is in the interpixel boundary if and only if x and y lie in different regions r_i and r_j , i.e. the grid line between x and y intersects the boundary B in at least one point v . The endcrack edgels are exactly the end points of these pixel edges, and their distance to v is at most $h/\sqrt{2}$. It follows that $q = h/\sqrt{2}$ for the endcrack digitization. The midcrack edgels are the center points between x and y , so their maximum distance to v is $\frac{h}{2}$. Hence, $q = \frac{h}{2}$ for the midcrack digitization. The maximum distance between neighboring edgels is h in both cases.

Now, we prove the bound on p given q . By definition $B = \bigcup \partial r_i$, where ∂r_i is the boundary of region r_i . Since every region contains a closed disc of radius $g = h/\sqrt{2}$, and every such disc contains at least one grid point, every region r_i contains a grid point, i.e. \hat{r}_i is not empty, and there exist at least four edgels near ∂r_i . Due to the nonexistence of (θ, d) -spikes, any two components $(\partial r_i)_j$ and $(\partial r_i)_l$ of the boundary ∂r_i must have a distance of more than $d \geq 4q$. So, for every component there exists a set of edgels which are closer to $(\partial r_i)_j$ than to any other component. Obviously every component $(\partial r_i)_j$ is a closed curve. Thus by mapping every edgel to the nearest point of B , one gets a cyclic list of points $[b_k]^{(ij)}$ for every component $(\partial r_i)_j$, and each point b_k has a distance of at most $h + 2q$ to its successor b_{k+1} in the list. For endcrack edgels, we have $h + 2q = (1 + \sqrt{2})h \leq d$, and for midcrack edgels $h + 2q = 2h \leq d$. Thus, the boundary part between b_k and b_{k+1} includes no point with an angle smaller than θ . As shown in Fig. 6, this implies that the distance from any boundary point between b_k and b_{k+1} to the nearer one of these two points is at most $(\frac{h}{2} + q) / \sin \frac{\theta}{2}$. Thus, the maximum distance to the nearest of the two edgels which are mapped onto b_k and b_{k+1} is $p = q + (\frac{h}{2} + q) / \sin \frac{\theta}{2}$. \square




Fig. 7. (a) Koch Snowflake; (b) subset digitization (note the topology violations); (c) (α, β) -boundary reconstruction from marked midcrack edgels. Areas where the edgels do not unambiguously determine the boundary shape pop out by remaining thick.

E.g. if a plane partition has no $(60^\circ, d)$ -spikes for sufficiently big d , we get $q = h/\sqrt{2}$ and $p = (1 + 3\sqrt{2})\frac{h}{4} \approx 1.31h$ for endcrack digitization and $q = \frac{h}{2}$ and $p = h$ for midcrack digitization. It follows that midcrack digitization should be favoured over endcrack digitization. The nonexistence of shape spikes allows us to topologically correctly digitize even objects having a fractal boundary like the Koch Snowflake (see Fig. 7): let K be the object bounded by the Koch Snowflake based on a triangle of sidelength 1. Then it can be shown that K is r -stable for all $r < 1/\sqrt{3}$ and it has no $(60^\circ, d)$ -spikes for $d < 1/\sqrt{3}$ and it contains a γ -disc for any $\gamma \leq 1/\sqrt{3}$. Thus the (α, β) -boundary reconstruction based on the midcrack digitization with a square grid of grid size h is correct for all $h < 1/\sqrt{27} \approx 0.192$.

Our sampling theorem can also be applied to commonly used edge detectors on real images (see Fig. 8), like Canny's algorithm [3]. In [12] we derive the following bounds: suppose the original partition is r -stable and free of $(60^\circ, 2r)$ -spikes, and the combined PSF and edge detector scale is at most $\sigma = 0.8r$, with pixel distance $h \leq r$. Then q does not exceed $0.9\sigma + 0.3 \approx 1.1$ pixels when the boundary contains corners or junctions and $\text{SNR} = 10$ (noise at this level is already quite visible), and $q \approx 0.2$ pixels when the partition is (4-pixel)-regular and $\text{SNR} = 30$. Note that the average error is much lower and approaches zero along straight edges. When the edgels are not represented with subpixel accuracy, a maximal round-off error of $\frac{h}{\sqrt{2}}$ must be added, and the average error cannot fall below $\frac{h}{\sqrt{6}} \approx 0.4h$ pixels even in case of straight edges.

5 Conclusions

To our knowledge, this paper proposes the first geometric sampling theorem that explicitly considers measurement errors. Moreover, the theorem applies to a much wider class of shapes (r -stable partitions) than existing ones (r -regular partitions). The situation in real images is thus modeled much more faithfully because shapes may now have corners and junctions, and standard segmentation algorithms can be used. The resulting segmentations are similar to what one gets from traditional heuristic edgel linking, but their properties can now be formally proven due to their theoretical basis in Delaunay triangulation. We




Fig. 8. Real images, (α, β) -boundary reconstruction (center) and minimal reconstruction after thinning (right). Edgels have been computed by Canny's algorithm on a color (top) and intensity (bottom) gradient.

showed that many known digitization and segmentation methods can be analysed and applied in the new framework by simply determining their error bounds. Our approach (including boundary thinning) provides a novel way for computing a combinatorial map representation [2] of the boundaries in real images.

References

1. Bernardini, F., Bajaj, C.L.: Sampling and Reconstructing Manifolds Using Alpha-Shapes. Proc. 9th Canadian Conf. Computational Geometry (1997)
2. Braquelaire, J.-P., Brun, L.: Image segmentation with topological maps and interpixel representation. J. Visual Communication and Image Repr. 9(1):62-79 (1998)
3. Canny, J.: A Computational Approach to Edge Detection. IEEE Trans. Pattern Analysis and Machine Intelligence, 8(6):679-698 (1986)
4. Edelsbrunner, H., Mücke, E.P.: Three-dimensional alpha shapes. ACM Trans. Graphics, 13:43-72 (1994)
5. Edelsbrunner, H.: The union of balls and its dual shape. Discrete Comput. Geom., 13:415-440 (1995)
6. Latecki, L.J., Conrad, C., Gross, A.: Preserving Topology by a Digitization Process. J. Mathematical Imaging and Vision, 8:131-159 (1998)
7. Pavlidis, T.: Algorithms for Graphics and Image Processing. Computer Science Press (1982)
8. Ronse, C., Tajine, M.: Discretization in Hausdorff Space. J. Mathematical Imaging and Vision 12:219-242 (2000)
9. Serra, J.: Image Analysis and Mathematical Morphology. Academic Press (1982)
10. Stelldinger, P., Köthe, U.: Towards a General Sampling Theory for Shape Preservation. Image and Vision Computing, 23(2):237-248 (2005)
11. Stelldinger, P.: Digitization of Non-regular Shapes. In: Ronse, C., Najman, L., Decenciere, E. (eds.): Mathematical Morphology, ISMM (2005)
12. Stelldinger, P., Köthe, U., Meine, H.: Topologically Correct Image Segmentation using Alpha Shapes. University Hamburg, Computer Science Department, Technical Report FBI-HH-M-336/06 (2006)