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 are not very realistic: they impose strong constraints on feasible shapes, and require measurements to be free of error. In this paper, we propose a new approach based on Delaunay triangulation and α-shapes which significantly weakens these restrictions. It assumes that sampling points (edgels) represent true object edges with a certain bounded error. We are able to prove under which conditions a topologically correct segmentation can be reconstructed from the edgels. Experiments on real and generated images demonstrate the good performance of the new method and confirm the predictions of our theory.

1 Introduction

The question, whether or when a computed image segmentation corresponds closely to the underlying real-world partitioning, is fundamental to image understanding. A number of partial results have been obtained in the past, but they are not sufficiently realistic to model many actual imaging situations. In this paper, we improve on the state-of-the-art by explicitly allowing measurement errors.

The analysis we are going to present is based on a clear distinction between the ideal geometric image, which cannot be observed in practice, and the actually available digital image. The geometric image has infinite resolution (i.e. is an analog function) and can be thought of as the geometric projection of a 3-dimensional scene, although we do not consider the details of the projection in this work. Instead, we think of the analog image as a given geometric partitioning of the plane into distinct regions. The interior of each region is described by some simple function (e.g. a constant), but the transitions between regions are discontinuous. This ideal analog image is then transformed into a digital image by a real camera. Beyond geometric projection, a real camera is characterized by its point spread function, the sampling grid and its quantization and noise models. Consequently, the partition of the geometric image must be inferred from the limited information in the digital image. We ask how accurate this reconstruction can be.

So far, geometric sampling theorems have been developed that describe the situation in 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 (definition below) and the sampling grid has a maximum pixel radius of at most $r' \leq r$ [11, 15]. By making slightly stronger assumptions ($r' + p \leq r$), this property is preserved when the shapes are blurred by a disc or square of radius p prior to discretization [9, 16]. It is even possible to relax the requirement of r-regularity somewhat into r-halfregularity [17], when the size of the regions is large enough.

However, the approaches mentioned have two important limitations. First, they are not applicable when there are more regions than just fore- and background. Second, they do not make any predictions of what happens when the segmentation can only be computed with a certain measurement error. One reason for these limitations is that the theorems are based on the assumption of a fixed (although sometimes arbitrary) sampling grid. In this paper, we are going to drop this assumption in favour of *adaptive sampling* where the sampling points are placed roughly along the contour of the regions to be segmented.

We can get to adaptive sampling in two ways: first, we can choose sampling points purely based on a fixed grid. This can be done by directly taking the grid points as sampling points, but drop all sampling points that are not near the segment boundaries. Thus we obtain a set of pixels marking segment edges (Fig. 4). Alternatively, we can choose points in the crack edge (inter-pixel boundary) between digital regions, i.e. on the pixel's dual grid (Fig. 5).

Second, we can allow sampling points to be placed freely in the plane. These points may, for example, result from the sub-pixel accurate version of Canny's algorithm or from exact contour following in a smoothly interpolated image by means of the predictor-corrector method or the sub-pixel watershed transform [10] (Fig. 8).

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* [6, 7]. The α-shape is essentially defined as the subset of the Delaunay triangulation of the points where the Delaunay cells' radius is below $\alpha \in \mathbb{R}^+$. 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. This implies that the regions are morphologically open and closed with respect to discs of radius $\leq r$, and that the curvature of the boundary cannot exceed $1/r$. This, in turn, means that regions cannot have corners, and that it is impossible to form junctions of three or more regions. 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 (connected components) must not change under either morphological opening or closing with a disc of radius $\leq r$. 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 called* r-stable *when its boundary* B *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 boundary point* $b \in B$ *to the nearest point in* S *is at most* p*, and the distance of every sampling 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 sampling points are exactly on the boundary, i.e.* $q = 0$*.*

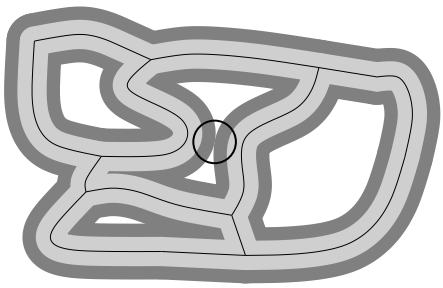

Fig. 1. An r-stable plane partition does not change the homotopy type when dilated with a disc of radius of at most r (light gray), while dilations with bigger radius (dark gray) may connect different arcs at waists as marked by the circle.

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, i.e. $d_H(S, B) = p$. Non-zero edgel shifts $q > 0$ can be caused by systematic or statistical measurement errors.

Edgels may be determined in various ways. For example, we can still use a fixed grid, but consider only those grid points which are near the boundary according to the boundary's grid intersection or supercover digitization (Fig. 4b). When a 4-connected region label image is given, we can define edgels as the nodes of the interpixel edge (crack edge). In this case, the edgels are a subset of the original grid's dual grid (Fig. 5). Often, it is possible to improve edge localization by shifting grid-based edgels to a sub-pixel position that is supposedly closer to the true boundary, as in Canny's algorithm (Fig. 8a). Edgels can also be defined by means of an edge-template matching procedure that provides sub-pixel locations in the first place, or be derived from a sub-pixel accurate edge tracing algorithm [10] (Fig. 8b).

For our discussion, the method used to compute edgels 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. This digital boundary does not need to be everywhere thin in the first step, but may include also some triangles instead of only edges. The idea o allow thick boundary representations is a major difference to previous edgel linking methods. It follows from the observation, that in reality the exact boundary can not be localized with arbitrary precision due to the erroneous digitization process. This also motivates the definition of r-stable plane partitions, where a thick representation of the boundary (i.e. a dilation with a disc) has the same homotopy type as the originally thin boundary. Such a thick boundary describes the area which includes the distorted reconstructed boundary. To make the process of boundary reconstruction by edgel linking more precise, we first need the Delaunay triangulation of a point set:

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 [6]:

Definition 5. *The α-complex D_α of a set of points S is defined as the subcomplex of the Delaunay triangulation D of S which contains all cells c such that*

- *the radius of the smallest open circumcircle of c is smaller than α, and this circle contains no point of S, or*
- *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 of the original plane partition we are analyzing. In order to determine when this is the case, the following theorem is of fundamental importance (the proof can be found in [7]):

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

Consequently, the polytope $|D_\alpha|$ 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, as the following theorem shows:

Theorem 2 (Bernardini & Bajaj). *Suppose the plane partition is r-regular, and S is a strict sampling of its boundary B such that $p \leq r$, $q = 0$. Then the polytope $|D_\alpha|$ is homotopy equivalent and even homeomorphic to the boundary B for all $p < \alpha < r$.*

This theorem is proved in [2]. Under these particular conditions, D_α does not contain any triangles – it only consists of edges and points and thus defines a plane partition in itself. According to the theorem, this plane partition is topologically equivalent to the original partition whose boundary was sampled by S. In other words, the α-complex completely defines the correct linking of edgels into edge chains. 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. 3 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.

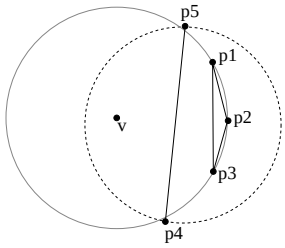

Fig. 2. Any circumcircle around p_4 and p_5 contains p_1, p_2, and p_3 (see text).

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. This is facilitated by the following definition:

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. The following lemma establishes that 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. Therefore, the radius of the triangle's circumcircle 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 the circumcircle of t because otherwise t would not be a Delaunay triangle. Neither t' nor e can contain v because their circumcircle radius would then be at least β. Now consider Fig. 2. It shows triangle t with corners p_1,

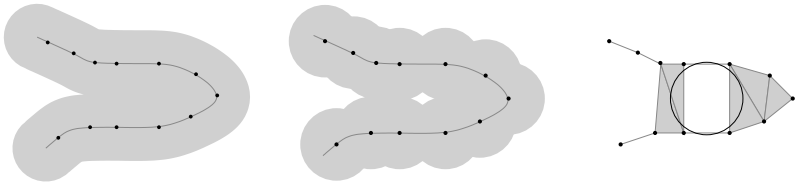

(a) (b) (c)

Fig. 3. It may happen that the α-dilation (a) of the boundary of an α-stable plane partition is not 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. These holes consist 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).

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 that it must not contain any other edgel. (In case of t', this condition follows from the fact that t' is a Delaunay triangle. In case of e, it stems from the definition of the α-complex.) The claim follows from the contradiction.

The reason for defining (α, β)-holes is that 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, as can be seen in Fig. 3. We solve this problem by introducing a second parameter β for the size of such holes, i.e. we use the notion of (α, β)-holes to "repair" α-complexes that contain too many holes:

Definition 7. *An (α, β)-boundary reconstruction from a set of edgels 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 edgel linking by hystheresis thresholding on the triangle size of a Delaunay triangulation:

1. Compute the Delaunay triangulation D of the edgels S.
2. Mark all triangles in D (including their edges) with a circumradius $< \alpha$.
3. Additionally mark Delaunay edges whose circumcircle contains no edgel and has a radius smaller than α.
4. Find connected components of unmarked triangles and edges.
5. For each component from step 4 which does *not* contain any triangle with a circumradius of at least β, mark all its triangles and edges.

The following theorem shows that exactly the desired holes survive when α and β are properly chosen. This theorem can be interpreted as a new sampling theorem for region boundaries.

Theorem 3 (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 $\mathcal{P}$, provided the following conditions are met:*

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 $B^{\oplus} = B \oplus \mathcal{B}^o_{\alpha+q}$ 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 the partition $\mathcal{P}$. Conversely, since $\alpha > p$, every open α-disc around a point of S intersects B, and the union U of these discs covers the entire boundary 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 $U \subset B^{\oplus}$ and $|D_\alpha| \subset U$. 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.

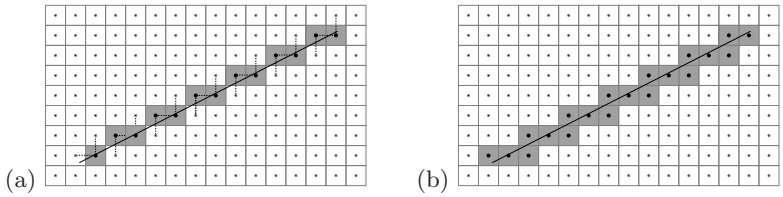

Fig. 4. (a) Where the boundary intersects the dual grid, the nearest sampling points form the *grid intersection digitization*. (b) The *supercover digitization* contains all sampling points whose pixel facets intersect the arc.

– 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 $|\mathcal{R}|$. □

If there exists no r such that all conditions of theorem 3 are fulfilled for a given plane partition (or if the chosen α is too big), it cannot be guaranteed that the regions of the (α, β)-boundary reconstruction have the same topology as the original ones. E.g. if a region is too small it may happen that it is lost in the reconstruction. If a region has an s-waist for $s \leq 2\alpha$, i.e. if the s-dilation of the boundary cuts the region into two or more components, it may happen that the region is also split into two or more parts in the (α, β)-boundary reconstruction. In case of very small waists, i.e. when $s + 2p + 2q \leq \alpha$, it is even guaranteed that the two waist sides are connected by at least one line segment in the reconstruction (and if the different parts of the original region are big enough, every part occurs in the reconstruction as a separate region). Thus, we can still apply our sampling theorem when a plane partition has some waists: we simply construct a new plane partition where the different sides of a waist are connected by a new arc. When the new partition fulfills the requirements, the modified topology is preserved, and the difference between the modified reconstruction and the original plane partition is well defined. The example shown in the middle row of Fig. 13 illustrates this fact.

4 Application to Popular Sampling and Segmentation Schemes

The parameters p and q in theorem 3 are assumed given. To make their meaning and consequences more intuitive, we compute or estimate these numbers for common sampling and segmentation schemes. Let's first look at grid intersection digitization:

Definition 8. *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*

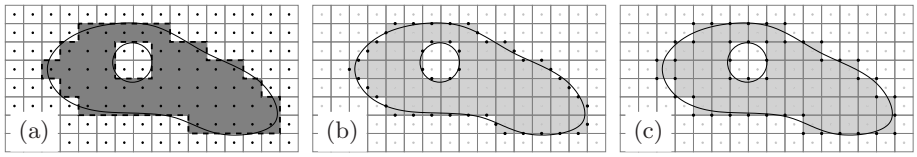

Fig. 5. The *interpixel boundary* (dashed) can be extracted from the subset digitization (a). It includes both the *midcrack digitization* (b) and the *endcrack digitization* (c).

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 3, 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 9. *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 the boundary (smaller spacing of the edgels) in the former. As stated in [14], the supercover digitization is a Hausdorff discretization, i.e. a set of sampling points which minimizes the Hausdorff distance to the boundary B. Since this Hausdorff distance is equal to $\max(p, q)$, 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 methods, in particular the subset digitization:

Definition 10. *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 the* midcrack digitization *(Fig. 5b) where the center points of all pixel edges inside the interpixel boundary $\mathbb{B}$ are chosen as edgels, and the* endcrack digitization *(Fig. 5c) where all pixel corner points lying on the interpixel boundary $\mathbb{B}$ 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, as the following considerations illustrate:

> 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 11. *Let $\mathcal{P}$ be a plane partition with boundary B. We say two points $x1, x2 \in \mathcal{B}$* 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. But this intuitive description cannot be used for the definition, since we want to apply it to fractal arcs as well. A fractal arc has infinite length, so the notion of shortest path is not applicable, but the arc may nevertheless be free of θ-spikes (see below).

It can be shown that r-regular partitions have no (θ, d)-spikes for all $d \leq r$ and $\theta = 2 \arctan\left(\frac{d}{2r - \sqrt{4r^2 - d^2}}\right)$ (e.g. for $\theta = 90°, 60°$ we get $d = r$ and $d = \sqrt{3}r$ respectively). By sampling the boundary of an r-regular partition dense enough, one can enforce the angles $\angle x_1 y x_2$ to be arbitrarily flat. In general, absence of (θ, d)-spikes does not imply r-stability, so we will require both. This property does not only restrict the maximal angle of the original contour. It also restricts the angle of any pair of adjacent line segments of the boundary reconstruction if the edgels on the boundary exceed a certain density. Fig. 6 shows that the distance from any point y on the path Y to the nearest point of x_1 and x_2 is at most $\frac{d}{2 \sin \frac{\theta}{2}}$. Although this restricts the angles of the reconstruction only for

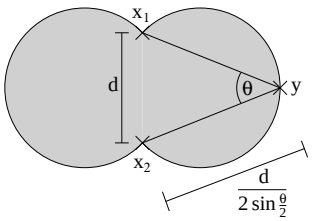

Fig. 6. Any point which encloses an angle of at most θ with x_1 and x_2 must lie inside the shaded region. The shown y is the one with the maximal distance to the nearer one of x_1 and x_2. Thus there is a path from x_1 to x_2 inside the shaded region and each of its points has a distance of at most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

noise-free sampling points, i.e. $q = 0$, it can be used for computing an upper bound for p even for digitizations with noise:

Theorem 4. *Let G be a square grid with sample distance h (pixel radius $g = \frac{h}{\sqrt{2}}$). Further, 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 of B is a (p, q)-boundary sampling with $q = \frac{h}{\sqrt{2}}$ and $p = q + \left(\frac{h}{2} + q\right) / \sin\frac{\theta}{2}$, provided that $h \leq \frac{d}{1+\sqrt{2}}$. Likewise, the midcrack digitization is a (p, q)-boundary sampling with $q = \frac{h}{2}$ and $p = q + \left(\frac{h}{2} + q\right) / \sin\frac{\theta}{2}$, 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 $\frac{h}{\sqrt{2}}$. It follows that $q = \frac{h}{\sqrt{2}}$ 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 on the interpixel boundary 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 $g = \frac{h}{\sqrt{2}}$, 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

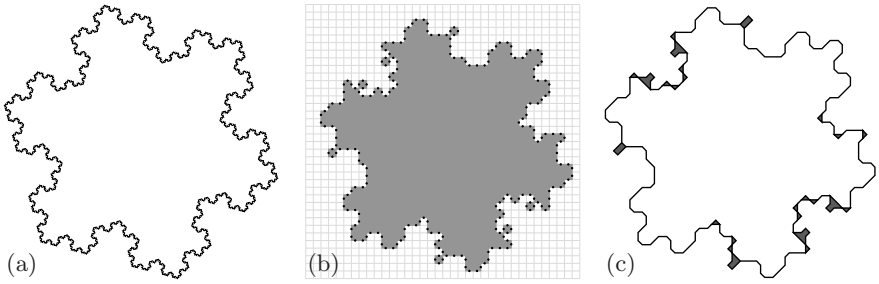

Fig. 7. (a) Koch Snowflake; (b) subset digitization of (a) with midcrack edgels marked (note the topology violations); (c) (α, β)-boundary reconstruction from midcrack edgels. Areas where the edgels do not unambigously determine the boundary shape pop out by remaining thick.

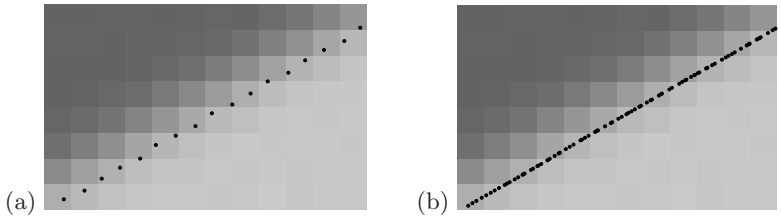

Fig. 8. Subpixel-accurate edgels from Canny's algorithm (a) and the subpixel watershed algorithm (b). Note the lower density and higher displacement of the former.

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 $\left(\frac{h}{2} + q\right) / \sin\frac{\theta}{2}$. Thus, the maximum distance to the nearest of the two edgels which are mapped onto b_k and b_{k+1} is $p = q + \left(\frac{h}{2} + q\right) / \sin\frac{\theta}{2}$. □

Thus for example, when $h = 1$ and the plane partition has no $(60°, d)$-spikes with $d > 2.4$, we get $p \approx 1.31$, $q \approx 0.7$ for endcrack and $p = 1$, $q = 0.5$ 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 $r < \frac{1}{\sqrt{3}}$ and it has no $(60°, d)$-spikes for $d < \frac{1}{\sqrt{3}}$ and it contains a γ-disc for any $\gamma \leq \frac{1}{\sqrt{3}}$. Thus the (α, β)-boundary reconstruction based on the midcrack digitization with a square grid of grid size h is correct for all $h < \frac{1}{\sqrt{27}} \approx 0.192$.

Many segmentation algorithms (e.g. zero-crossing-based edge detectors and the watershed algorithm) compute image labelings similar to subset digitization, which can be used to define endcrack and midcrack edgels. However, their error bounds differ from the ideal ones obtained above. To quantify these differences,

we model the transformation from analog to digital images in real cameras:

$$f_{ij} = (\mathrm{PSF} \star f(x,y))_{ij} + n_{ij} \tag{1}$$

where $f(x,y)$ is the ideal geometric image, PSF is the point spread function, subscripts denote sampling, and n_{ij} is additive Gaussian noise (quantization is neglected). The PSF (which shall be band-limited) suppresses high spatial frequencies and the resulting smooth transitions between regions allow for subpixel accurate edge localization. On the other hand, systematic localization errors are introduced because blurring distorts edges. Noise causes additional statistical errors in p and q. We estimate these errors for a number of exemplary edge detectors: we consider two variants of the Haralick detector as representatives of zero-crossing-based algorithms, and three variants of Canny's algorithm to exemplify ridge-based edge detection. Haralick [8] defines edgels at the zero-crossing of the second derivative along the gradient direction:

$$b = f_x^2 f_{xx} + 2 f_x f_y f_{xy} + f_y^2 f_{yy} \stackrel{!}{=} 0 \tag{2}$$

provided that the third derivative along the same direction is negative (indicating a local gradient maximum), and the gradient magnitude is above a threshold. Crack edges between positive and negative pixels of b where the constraints are fulfilled define a set of midcrack edgels. Their fixed accuracy can be improved when a continuous function $\tilde{b}$ is computed by spline interpolation of b, and edgels are located in $\tilde{b}$ by means of Newton iteration along the gradient direction. In our implementation of this variant, edgels are placed roughly at a distance of 0.1 pixels along the edge, Fig. 9a, b.

In contrast, Canny's algorithm [5] uses the gradient magnitude $\sqrt{f_x^2 + f_y^2}$ and looks for relative maxima along the gradient direction. Better localization (significantly smaller q) is achieved by either computing the maximum of an approximating parabola accross the edge, or by Newton iterations on a continously interpolated version of the gradient image, Fig. 9c and d. We estimate p and q on a large number of images created by numerical solution of the convolution integral (1) at various angles and grid positions, Fig. 9. Derivatives are computed by Gaussian filters at scale σ_E, and the PSF is also Gaussian with scale σ_{PSF}. To avoid aliasing we use $\sigma_E \geq 1$ and $\sigma_{\mathrm{PSF}} = 1$ (cf. [18]).

First, consider noise-free straight edges. A radially symmetric PSF does not distort straight edges and q should be close to zero (non-zero values reflect discrepancies between the computational theory and its actual realization). Subpixel methods achieve $q \lesssim 0.05$ pixels. With the exception of the subpixel Haralick operator (which places edgels very densely), p roughly equals the pixel radius. Row 1 in Table 1 lists the maximum errors we found.

The effect of image noise on straight edge localization was analysed by Canny [5]. When the noise is Gaussian distributed with zero mean and standard deviation s_N, the expected error (in pixels) is

$$E[\xi] = \frac{s_N}{a} \frac{\sqrt{6}}{4} \left(1 + \frac{\sigma_{PSF}^2}{\sigma_E^2}\right)^{3/2} \tag{3}$$

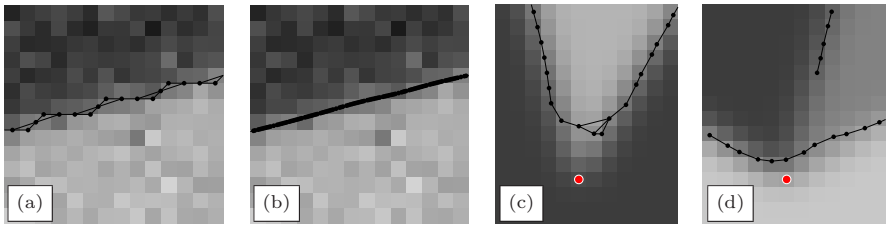

Fig. 9. Edgels and boundary reconstruction using $\alpha = 1.55, \beta = 2$: (a) midcrack variant and (b) subpixel variant of Haralicks algorithm. Note the lower density and higher displacement of the former. (c) Parabola and (d) spline variant of Canny's algorithm. Red dots indicate the ground-truth corner locations.

where a is the height of the step, and a/s_N is the signal-to-noise ratio (SNR). When $\sigma_{PSF} \approx \sigma_E$, we get $E[\xi] \approx 1.7\frac{s_N}{a}$. For $\sigma_E \to \infty$, the error approaches $0.6\frac{s_N}{a}$ (the common belief that the error increases with σ_E is only justified in 1D). In typical images $\frac{a}{s_N}$ is between 5 and 100. The expected statistical error is then below 0.2 pixels, and the maximum error does not exceed $3E[\xi] = 0.6$ pixels with probability 0.997. Rows 2 and 3 of Table 1 confirm these predictions.

Smoothing of curved boundaries with the PSF results in biased edgel positions. The gradient magnitude of a disc with radius ρ and contrast a is [3]

$$g(r) = |a|\frac{\rho}{\sigma^2}e^{-\frac{r^2+\rho^2}{2\sigma^2}}I_1\left(\frac{r\rho}{\sigma^2}\right) \tag{4}$$

where r is the distance from the center of the disc, I_1 is the modified Bessel function of order 1, and $\sigma^2 = \sigma_{PSF}^2 + \sigma_e^2$ is the combined scale of the PSF and edge operator. The bias depends on the curvature radius ρ and the scale σ. It is directed towards the concave side of the curve when $\sigma < 0.8\rho$ (which is true in most practical situations). Row 4 of Table 1 compares theoretical predictions and experimental estimates for $\rho = 4$. It can be seen that the best methods (using spline interpolation and Newton iterations) are very close to the theoretical limit.

A bias toward the concave side of the contour is also observed at corners. Its magnitude depends on σ and the corner angle φ and is maximal along the bisector of the corner. The gradient maximum along the bisector (i.e. the estimated edge location) is the solution of the implicit equation [13]

$$\frac{1}{2\pi\sigma^2}e^{-\frac{r^2}{2\sigma^2}} - \left(\tan\left(\frac{\varphi}{2}\right)\right)^2\frac{r}{2}\left(1+\operatorname{erf}\left(\frac{r}{\sqrt{2}\sigma}\right)\right) = 0 \tag{5}$$

where erf is the error function. The sharper the corner, the higher the bias. E.g. for $\varphi = 90°, 45°, 15°$ it is approximately 0.5σ, 1.2σ, and 2.2σ. Rows 5 and 6 in Table 1 show that actual errors are even higher than theory predicts.

The situation at junctions is even more complicated. The large number of degrees of freedom (angles, intensities) does not allow the error to be described in a compact way. The algorithms considered here are usually unable to close all contours near a junction. The remaining gaps also cause p to attain quite large values, as row 7 of Table 1 shows.

Table 1. Experimental estimates of the maximum errors p and q (pixels). Theoretical predictions are given in brackets. Unless noted, there was no noise and $\sigma_{\mathrm{PSF}} = \sigma_E = 1$.

	Canny (pixel coordinates)		Canny (parabola)		Canny (spline)		Haralick (midcrack)		Haralick (spline)	
	p	q	p	q	p	q	p	q	p	q
straight line	0.79	0.70 [0.7]	0.71	0.05 [0.0]	0.75	0.02 [0.0]	0.70	0.47 [0.5]	0.19	0.46 [0.0]
straight line SNR = 10	1.0	0.82	0.81	0.47 [0.52]	0.92	0.57 [0.52]	0.90	0.93	0.63	0.85 [0.52]
straight line $\sigma_E = 2$, SNR = 10	1.0	0.81	1.0	0.28 [0.26]	1.0	0.28 [0.26]	0.79	0.73	0.57	0.81 [0.26]
disc, radius = 4		0.73		0.73		0.25 [0.2]		0.74		0.29 [0.2]
corner, 90°	1.58 [0.71]	0.84	1.38 [0.71]	0.76	1.34 [0.71]	0.69	1.52 [0.71]	0.93	1.15 [0.71]	0.71
corner, 15°	4.03 [3.1]	1.3	3.99 [3.1]	0.92	3.96 [3.1]	0.94	3.39 [3.1]	1.33	3.96 [3.1]	1.3
junction, degree = 3	2.70	1.56	2.66	1.15	2.70	1.40	2.25	1.81	2.20	1.71

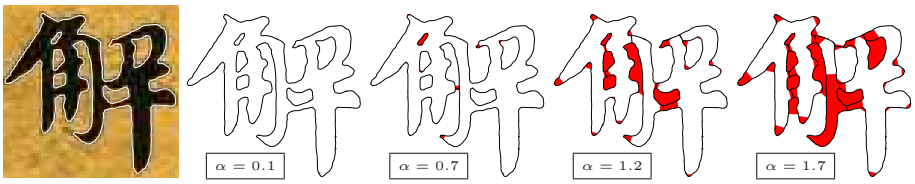

Fig. 10. Chinese character (*white*: contours extracted by levelcontour tracing [1]), (α, β)-boundary reconstructions with increasing values of α (*red*: before thinning, *black*: minimal boundary reconstruction)

Fig. 13 and Fig. 10 show results of α, β-reconstruction in real images. Region topology is correctly recovered when α and β are properly chosen. Since edgels are considered as isolated points, our new algorithm also facilitates the combination of edgels from different sources, cf. Fig. 11: The edgels computed by Canny's algorithm are not very accurate near corners and junctions, and this requires large α and β causing the reconstruction to be thick in problematic areas (gray). In a second step, a maximum likelihood junction position is computed from the gradient magnitudes and directions at the edgels in a neighborhood of each thick area, resulting in the red points. These points are simply added to the set of edgels, and the reconstruction from the new set is much more accurate than the original one.

Taking everything together, we arrive at the following approximate bounds: suppose the original partition is r-stable and free of $(60°, 2r)$-spikes (i.e. corners enclose at least 60°, curved arcs have at least curvature radius $\rho = 2r/\sqrt{3}$), and the combined PSF and edge detector scale is at most $\sigma = 0.8r$. Moreover, σ should not be smaller than 0.9 pixels in order to avoid aliasing [18], so the pixel distance must be $h \approx r$. Then q does not exceed $0.9\sigma + 0.3 \approx 1.1$ pixels when the boundary contains corners or junctions and SNR = 10 (this is quite visible noise), and $q \approx 0.2$ pixels when the partition is (4-pixel)-regular and SNR = 30. Note that these bounds are maximum errors, the average error is much lower and approaches zero along straight edges. When the edgels are not

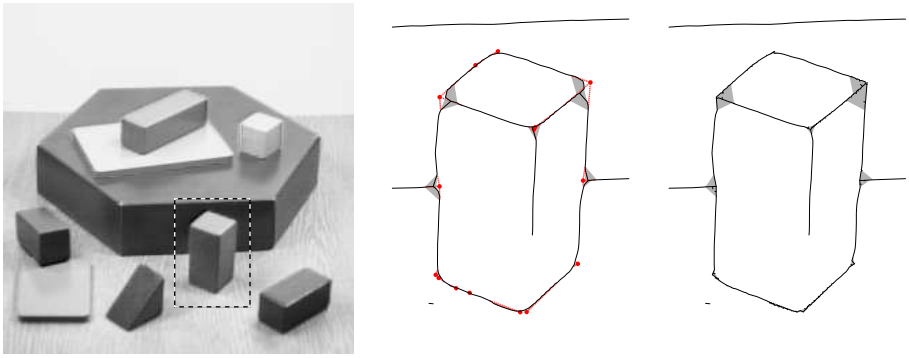

Fig. 11. Left: original image and ROI; center: (α, β)-boundary reconstruction from subpixel Canny edgels (black and gray), thinned reconstruction (black only) and additional edgels to be added (red); right: modified reconstruction including new edgels.

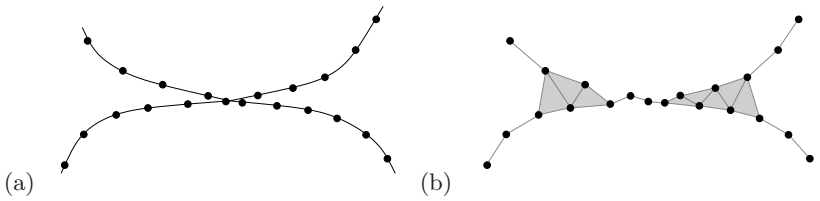

Fig. 12. Narrow spikes can lead to a boundary reconstruction where originally unconnected regions (a) look like they had a common boundary edge (b).

represented with subpixel accuracy, a round-off error of $h/\sqrt{2}$ must be added, and the average error cannot fall below 0.4 pixels (the standard deviation of a uniform distribution in the unit square) even in case of straight edges.

5 Boundary Thinning and Neighborhood Relations

The sampling theorem presented above tells us how to reconstruct all regions of a plane partition with correct topology and how to reconstruct the boundary of the partition with correct homotopy type. But it does not say anything about the preservation of neighborhood relations, i.e. whether the reconstructions of two regions have a common boundary if the original regions did have one. First, it is not straightforward to define neighborhood when the boundary representation can be thick (i.e. may contain triangles). Second, two regions whose reconstructions are connected (i.e. have a common thin boundary), have not necessarily been connected in the original partition, as can be seen in Fig. 12.

Many algorithms that build upon segmentation results cannot handle partially thick boundary representations. We can recover a thin (i.e. locally 1-dimensional) boundary by *topology-preserving thinning*, which works similar to

skeletonization in a pixel-based region representation. 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 can be easily recognized: they bound an (α, β)-hole on one side and a triangle in the boundary reconstruction on the other. Thinning removes simple edges until any further removal would change the topology (i.e. create an isolated sampling point or merge two regions). The algorithm is as follows:

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: get the topmost edge from the queue and check whether it is still simple (it may have lost this property after removal of other edges). If yes, remove it from the boundary reconstruction. Check whether the other edges of the same triangle have now become simple and put them in the queue if this is the case.

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) along each edge and remove weak edges first. A particularly interesting ordering is defined by the length of the edges:

Definition 12. *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. 13. Since region topology is preserved, the minimal boundary reconstruction is homotopy equivalent to the boundary B of the original plane partition ∂P. The two boundaries do not in general have the same topology, because the adjacency relations between regions may differ (see below for details), and the reconstruction may contain "dangling" edges, which end in the interior of a region. Since only the shortest edges survive, dangling edges cannot reach very far into a region and don't pose problems.

Since a minimal boundary reconstruction can be shown to be a shortest possible one with correct topology, the surviving edges connect edgels closest to each other. Neihgboring edgels therefore align in an optimal way 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 between a true boundary point and the nearest edgel may be much larger than $d_{\max}/2$ if the displacement of neighboring edgels is highly correlated as is usually the case in practice. For example, edgels along a circular arc are consistently biased toward the concave side of the curve. When we set $\alpha' = d_{\max}/2 + \epsilon < p$ (with arbitrarily small ϵ), an (α', β) reconstruction of the edgel set is still correct in the sense of theorem 3: 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 remained 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.

Since a minimal boundary reconstruction does not necessarily reproduce the original region adjacencies, it is interesting to ask whether some neighborhood relations can nevertheless be recovered from a (p, q)-boundary sampling and the associated (α, β)-boundary reconstruction. We showed in Fig. 12 that false adjacencies can be caused by narrow spikes in the original boundaries. When these boundaries are free of (θ, d)-spikes, thin parts of the boundary reconstruction that exceed a certain length can never arise from unfortunate spike configurations, but reflect the true adjacency of two original boundaries:

Theorem 5. *Let $\mathcal{P}$ be an r-stable plane partition with regions r_i and boundary B having no (θ, d)-spikes. Further, let S be a (p, q)-sampling of B and $\mathcal{R}$ the (α, β)-boundary reconstruction of S with regions h_i, such that all requirements of theorem 3 are fulfilled. $S_i = \partial h_i \cap S$ denotes the set of edgels on the boundary of h_i. When $d \geq 2(\alpha + q)$ and $p' := d/\left(2 \sin \frac{\theta}{2}\right) + q$ the following holds:*

1. *If the distance between the two nearest edgels of S_i and S_j exceeds $2p'$, the corresponding original regions r_i, r_j are not adjacent, i.e. $\partial r_i \cap \partial r_j = \emptyset$.*
2. *When there exists a point x with $d_H(x, S_i) \leq p'$, $d_H(x, S_j) \leq p'$ and $d_H(x, S_k) > 2p'$ for all $k \neq i, j$, the original regions r_i, r_j are arc-adjacent.*
3. *If two regions r_i, r_j have a distance greater than $2(p' + q)$, the conditions of item 1 are always fulfilled.*
4. *If two regions r_i, r_j have a common boundary point x such that $d_H(x, r_k) > 3p'$ for all $k \neq i, j$, the conditions of item 2 are always fulfilled, i.e. adjacency of r_i and r_j can be detected in the boundary reconstruction.*

Proof. (1) For any $s_t \in S_i$ let $x_t \in \partial r_i$ be the nearest boundary point. Then for any two s_{t_1}, s_{t_2} being connected by a line segment of ∂h_i, the distance between x_{t_1} and x_{t_2} is smaller than $2(\alpha + q)$. Since (θ, d)-spikes do not exist, the distance of each point of ∂r_i to the nearest x_t cannot exceed $d/\left(2 \sin \frac{\theta}{2}\right)$ and thus the distance of ∂r_i to ∂h_i is bounded by p'. The same holds for h_j. When the shortest distance between S_i and S_j is larger than $2p'$, ∂r_i and ∂r_j cannot intersect.
(2) Both S_i and S_j intersect the disc $\mathcal{B}^0_{p'}(x)$. Since $d_H(x, S_k) > 2p'$ for every $k \neq i, j$, no part of ∂r_k can intersect $\mathcal{B}^0_{p'}(x)$. Thus r_i and r_j are the only regions which intersect $\mathcal{B}^0_{p'}(x)$, which is only possible when they have a common edge.
(3) Since the distance between r_i and r_j exceeds $2(p' + q)$, S_i, S_j have to be more than $2p'$ away from each other.
(4) Due to the absence of (θ, d)-spikes, the distance $d_H(x, S_k), k \neq i, j$ must be greater than $2p'$. For the same reasons, $d_H(x, S_i) \leq p'$ and $d_H(x, S_j) \leq p'$. □

If every junction of $\mathcal{P}$ has degree 3, the boundary sampling only needs to be sufficiently dense (i.e. α has to be sufficiently small) in order to reconstruct

not only the topology of every region of a plane partition, but also the complete neighborhood relations, i.e. the complete combinatorial map [4] encoding $\mathcal{P}$'s abstract topology, without any error.

6 Conclusions

To our knowledge, this paper proposes the first geometric sampling theorem that explicitly considers measurement errors. Moreover, our new theorem applies to a much wider class of shapes (r-stable partitions) than existing theorems (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. We carefully derive the theoretical properties of several well-known edge detectors in order to apply our new theorem and demonstrate theoretically correct edgel linking. The resulting segmentations are similar to what one gets from traditional heuristic edgel linking, but their properties can now be formally proven thanks to their theoretical basis in Delaunay triangulation. The key to these advancements has been the shift of attention from region-based digitization models to edge based ones: the assumption that no sampling points are in the interior of any region (beyond the known error bound) allows us to reliably recover region and boundary connectivity. Our approach (including boundary thinning) provides a novel way for computing a combinatorial map representation [4] of the boundaries in real images.

We demonstrated that many known digitization and segmentation methods can be analyzed and applied in the new framework by simply determining their error bounds. We can predict whether a given image will be handled properly by an algorithm with a certain error bound. When the error increases, the performance degrades gracefully: first, the recovered boundary becomes thick when the detailed curve shape or junction connectivity can no longer be unambigously determined. Then, regions get split at too narrow waists, and finally too small regions will be lost. When additional edgels are added within the thick part of the (α, β)-boundary reconstruction, the accuracy parameters p and q will never increase. This opens up new possibilities for algorithm combination. For example, one could start with an edge detector, which produces thick boundaries near corners and junctions. Additional edgels can then be computed by a corner detector whose output is confined to these areas, so that it cannot produce false positives within regions. In future research we will investigate how false positives (large q) and false negatives (large p) can be recognized and removed.

References

1. E.L. Allgower, K. Georg: Numerical path following. In: P.G. Ciarlet, J.L. Lions (Eds.), Handbook of Numerical Analysis **5** (1997) 3–207. North-Holland.
2. F. Bernardini, C.L. Bajaj: *Sampling and Reconstructing Manifolds Using Alpha-Shapes*, Proc. 9$^{\text{th}}$ Canadian Conf. Computational Geometry, 1997

3. H. Bouma, A. Vilanova, L.J. van Vliet, and F.A. Gerritsen: *Correction for the dislocation of curved surfaces caused by the PSF in 2D and 3D CT images*, IEEE Trans. Pattern Analysis and Machine Intelligence, 27(9):1501-1507, 2005
4. J.-P. Braquelaire, L. Brun: *Image segmentation with topological maps and interpixel representation*, J. Visual Communication and Image Representation 9(1):62-79, 1998
5. J. Canny: *A Computational Approach to Edge Detection*, IEEE Trans. Pattern Analysis and Machine Intelligence, 8(6):679-698, 1986
6. H. Edelsbrunner, E.P. Mücke: *Three-dimensional alpha shapes*, ACM Trans. Graphics, 13:43-72, 1994
7. H. Edelsbrunner: *The union of balls and its dual shape*, Discrete Comput. Geom., 13:415-440, 1995
8. R. Haralick, L. Shapiro: *Computer and Robot Vision*, vol. 1, Addison Wesley, 1992
9. L.J. Latecki, C. Conrad, A. Gross: *Preserving Topology by a Digitization Process*, J. Mathematical Imaging and Vision, 8:131-159, 1998
10. H. Meine, U. Köthe: *Image Segmentation with the Exact Watershed Transform*, in: J.J. Villanueva (Ed.): VIIP 2005, Proc. 5$^{\text{th}}$ IASTED International Conference on Visualization, Imaging, and Image Processing, pp. 400-405, ACTA Press, 2005
11. T. Pavlidis: *Algorithms for Graphics and Image Processing*, Computer Science Press, 1982
12. L. Prasad, A.N. Skourikhine: *Vectorized Image Segmentation via Trixel Agglomeration*, in: L. Brun, M. Vento (Eds.): Graph-Based Representations in Pattern Recognition, Proc. 5$^{\text{th}}$ IAPR Workshop GbRPR '05, LNCS 3434, pp. 132-141, Springer, 2005
13. K. Rohr: *Localization Properties of Direct Corner Detectors*, J. Mathematical Imaging and Vision 4:139-150, 1994
14. C. Ronse, M. Tajine: *Discretization in Hausdorff Space*, J. Mathematical Imaging and Vision 12:219-242, 2000
15. J. Serra: *Image Analysis and Mathematical Morphology*, Academic Press, New York, 1982
16. P. Stelldinger, U. Köthe: *Towards a General Sampling Theory for Shape Preservation*, Image and Vision Computing, 23(2):237-248, 2005
17. P. Stelldinger: *Digitization of Non-regular Shapes*, in: C. Ronse, L. Najman, E. Decenciere (Eds.): Mathematical Morphology, Proc. 7$^{\text{th}}$ Intl. Symposium on Mathemat. Morphology, Springer, 2005
18. B. ter Haar Romeny: *Front-End Vision and Multi-Scale Image Analysis*, Kluwer Academic Publishers, Dordrecht, 2003

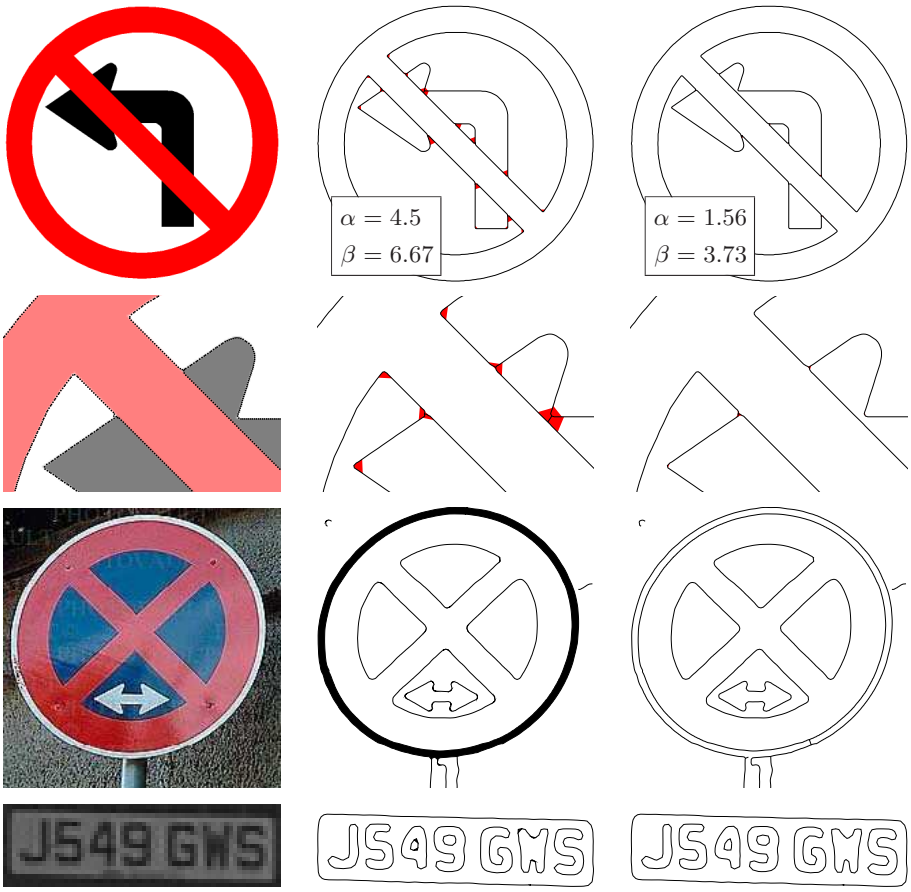

Fig. 13. *Top row*: generated image, reconstructions before (red and black) and after (black only) boundary thinning. *Second row*: details (left: original with edgels). Note the connectivity error in the center image where α is too big. *Third and fourth rows*: real images (left), (α, β)-boundary reconstruction (center) and minimal reconstruction after thinning (right).Edgels have been computed by Canny's algorithm on a color (third row) and intensity (fourth row) gradient.