

Provably Correct Edgel Linking and Subpixel Boundary Reconstruction

Ulrich Köthe, Peer Stelldinger, and Hans Meine

University of Hamburg, 22527 Hamburg, Germany

Abstract. Existing methods for segmentation by edgel linking are based on heuristics and give no guarantee for a topologically correct result. In this paper, we propose an edgel linking algorithm based on a new sampling theorem for shape digitization, which guarantees a topologically correct reconstruction of regions and boundaries if the edgels approximate true object edges with a known maximal error. 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, since they do not allow 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 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. 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.

Recently we developed a geometric sampling theorem which assumes that sampling points (edgels) are placed roughly along the contour of the regions to be segmented. The edgels can be obtained by an arbitrary edge detector, as long as the accuracy of the detected edges is known. In this paper, we compare common edge detectors in the context of our theory and show how to use them for generating a topologically correct image segmentation.

2 The Boundary Reconstruction Algorithm

We consider the task of reconstructing the boundary of a partition of the Euclidean plane from a sampled representation. The plane partition \mathcal{P} to be recovered is defined by a finite set of *points* $P = \{p_i \in \mathbb{R}^2\}$ and a set of pairwise disjoint *arcs* connecting these points. 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 .

Previous proofs about topologically correct reconstruction were restricted to *binary* partitions. That is, one 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. Examples are *r-regular partitions* in [5, 6, 8, 9] and *r-halfregular partitions* in [10]. Both are too restrictive for practical use (see [11] for details). In this paper we use a more general class of feasible plane partitions:

Definition 1. *A plane partition \mathcal{P} 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 $s \leq r$. 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 no pair of boundary points $x_1, x_2 \in B$ delimits a (θ, d) -spike.*

Thus a plane partition is r -stable if 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. Obviously, an r -stable plane partition has no $(\pi, 2r)$ -spikes. Intuitively, two points delimit a (θ, d) -spike, if the shortest boundary path between them does not differ too much from a straight line – it lies inside the shaded region in Fig. 2. In order to digitize such plane partitions, we approximate the *boundary* of the partition with a finite set of *adaptively placed* sampling points. The sampling points should be “near” the boundary:

Definition 2. *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 points in S are called edgels.*

The Hausdorff distance $d_H(B, S)$ between the boundary and the sampling points is $\max(p, q)$. The exact values of p and q depend on where the edgels come from. This is discussed in detail in section 4. Our new edgel linking algorithm is essentially a hysteresis thresholding on the sizes of Delaunay triangles:

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.




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 as marked by the circle (see Def. 1).




Fig. 2. Any point which encloses an angle of at least 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 $\frac{d}{2 \sin \frac{\theta}{2}}$.

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 union of marked triangles and edges is a simplicial complex which we denote (α, β) -*boundary reconstruction* from the edgels. The components of its complement are called (α, β) -*holes*. Under certain conditions, these holes exactly correspond to the regions of the original r -stable plane partition, as proven in [11]:

Theorem 1 (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 $\gamma \geq \beta + q > 2(p + q)$.

3 Boundary Thinning and Neighborhood Relations

Since the (α, β) -boundary reconstruction may contain triangles, it is not in general thin (i.e. locally 1-dimensional). However, many algorithms that build upon segmentation results cannot handle partially thick boundary representations. Therefore we propose a topology preserving boundary thinning. We call an edge in the (α, β) -boundary reconstruction *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 all simple edges iteratively:

1. Find all simple edges of the given (α, β) -boundary reconstruction and put them in a priority queue (the sorting is discussed below).




Fig. 3. *left:* original; *center:* (α, β) -boundary reconstruction; *right:* minimal reconstruction after thinning. (Edgels from Canny's algorithm on a color gradient)

2. As long as the queue is not empty, fetch an edge from the queue and remove it from the boundary reconstruction if 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) along each edge and remove weak edges first. A particularly interesting ordering is defined by the length of the edges:

Definition 3. *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, the minimal boundary reconstruction is homotopy equivalent to the boundary B of the original plane partition ∂P . The two boundaries are not in general topologically equivalent, because the adjacency relations between regions may differ (see below for details), and the reconstruction may contain short edges, which end in the interior of a region (they can also be removed iteratively).

Since the minimal boundary reconstruction is the shortest possible one with correct topology, the surviving edges connect edgels closest to each other. Neighboring 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 1: since the 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.




Fig. 4. Narrow spikes can lead to a boundary reconstruction where originally unconnected regions (a) look like they had a common boundary edge (b).

Theorem 1 does not guarantee that the neighborhood relations between reconstructed regions are the same as of the original regions, as can be seen in Fig. 4. The following theorem shows that neighborhood relations are preserved when the boundary arcs are long enough and free of (θ, d) -spikes:

Theorem 2. *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 1 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 $p' := d / (2 \sin \frac{\theta}{2}) + q$ 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 $d / (2 \sin \frac{\theta}{2})$ 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 $\mathcal{B}_{p'}^0(x)$. Since $d_H(x, S_k) > 2p'$ for every $k \neq i, j$, no part of ∂r_k can intersect $\mathcal{B}_{p'}^0(x)$. Thus r_i and r_j are the only regions which intersect $\mathcal{B}_{p'}^0(x)$, 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 $d_H(x, S_j) \leq p'$. \square

It follows that if every junction of \mathcal{P} has degree 3, the boundary sampling only needs to be sufficiently accurate (i.e. p, q , and α are 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. a complete combinatorial map [2] encoding \mathcal{P} 's abstract topology, without any error.




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).

4 Application to Popular Segmentation Schemes

In order to apply our boundary reconstruction algorithm, we can derive correct choices for α and β from the error bounds p and q of the edgel detector. First, let us pretend that we have access to the exact projected image, i.e. to the plane partition \mathcal{P} . One possibility to digitize this partition is the so-called *subset digitization*: We assign the same label to two pixels iff their centers are in the same region. Then, interpixel edges (crack edges) can be defined between pixel facets with different labels, see Fig. 5a. Crack edges give rise to two natural kinds of edgels: *endcrack* and *midcrack* edgels (located on the end or center points of the cracks respectively, Fig. 5b and c). When the boundaries of the plane partition are free of (θ, d) -spikes, the following bounds can be derived [11]: $q = \frac{h}{\sqrt{2}}$ (endcrack) and $q = \frac{h}{2}$ (midcrack) and $p = q + (\frac{h}{2} + q) / \sin \frac{\theta}{2}$ (both cases), where $h \leq \frac{d}{1+\sqrt{2}}$ is the required pixel distance. For example, when $h = 1$ and the plane partition has no $(60^\circ, d)$ -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, i.e. the latter is more accurate.

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:

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

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 sub-pixel 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 [4] defines edgels at the zero-crossing of the second derivative along the gradient direction:

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




Fig. 6. Edgels and boundary reconstruction using $\alpha = 1.55, \beta = 2$: (a) midcrack variant and (b) subpixel variant of Haralick's 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.

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. 6a, b.

In contrast, Canny's algorithm [3] 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 across the edge, or by Newton iterations on a continuously interpolated version of the gradient image, Fig. 6c 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. 6. 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 $\sigma_{PSF} = 1$ (cf. [12]).

First, consider straight edges. A radial 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 [3]. When the noise is Gaussian distributed with zero mean and standard deviation s_N , the expected error (in pixels) is

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

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 $\sigma_E \rightarrow \infty$, the error approaches $0.6 \frac{s_N}{a}$ (the common belief that the error increases with σ_E is only justified in

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 $\sigma_{\text{PSF}} = \sigma_E = 1$.

	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

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 [1]

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

where r is the distance from the center of the disc, I_1 is the modified Bessel function of order 1, and $\sigma^2 = \sigma_{\text{PSF}}^2 + \sigma_e^2$ 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 [7]

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

where erf is the error function. The sharper the corner, the higher the bias. E.g. for $\varphi = 90^\circ, 45^\circ, 15^\circ$ it is approximately $0.5\sigma, 1.2\sigma$, 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.




Fig. 7. Chinese character (*white*: contours extracted by levelcontour tracing [13]), (α, β) -boundary reconstructions with increasing values of α (*red*: before thinning, *black*: minimal boundary reconstruction)

Fig. 3 and Fig. 7 show results of α, β -reconstruction in two 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. 8: 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.

5 Conclusions

To our knowledge, this paper exploits the first geometric sampling theorem which explicitly considers measurement errors. We carefully derive the theoretical properties of several well-known edge detectors in order to apply our new theorem and demonstrate theoretically correct edge linking. The resulting segmentations are similar to what one gets from traditional heuristic 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.

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 unambiguously determined. Then, regions get split at too narrow waists, and finally too small regions will be lost (cf. Fig. 7). 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 and a large α which produces




Fig. 8. Left: original image and ROI; center: (α, β) -boundary reconstruction from sub-pixel Canny edges (black and gray), thinned reconstruction (black only) and additional edges to be added (red); right: modified reconstruction including new edges.

thick boundaries near corners and junctions. Additional edges 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 fact, false positives (large q) and false negatives (large p) are the major difficulties in our new algorithm. We are currently investigating how these can be recognized and removed.

References

1. 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
2. J.-P. Braquelaire, L. Brun: *Image segmentation with topological maps and interpixel representation*, J. Visual Communication and Image Repr. 9(1):62-79, 1998
3. J. Canny: *A Computational Approach to Edge Detection*, IEEE Trans. Pattern Analysis and Machine Intelligence, 8(6):679-698, 1986
4. R. Haralick, L. Shapiro: *Computer and Robot Vision*, vol. 1, Addison Wesley, 1992
5. L.J. Latecki, C. Conrad, A. Gross: *Preserving Topology by a Digitization Process*, J. Mathematical Imaging and Vision, 8:131-159, 1998
6. T. Pavlidis: *Algorithms for Graphics and Image Processing*, Comp. Sc. Press, 1982
7. K. Rohr: *Localization Properties of Direct Corner Detectors*, J. Mathematical Imaging and Vision 4:139-150, 1994
8. J. Serra: *Image Analysis and Mathematical Morphology*, Academic Press, 1982
9. P. Stelldinger, U. Köthe: *Towards a General Sampling Theory for Shape Preservation*, Image and Vision Computing, 23(2):237-248, 2005
10. P. Stelldinger: *Digitization of Non-regular Shapes*, in: C. Ronse, L. Najman, E. Decenciere (Eds.): *Mathematical Morphology*, ISMM, 2005
11. P. Stelldinger, U. Köthe, H. Meine: *Topologically Correct Image Segmentation using Alpha Shapes*, University Hamburg, Technical Report, 2006
12. B. ter Haar Romeny: *Front-End Vision and Multi-Scale Image Analysis*, Kluwer Academic Publishers, Dordrecht, 2003
13. Allgower, E.L., Georg, K.: Numerical path following. In: P.G. Ciarlet, J.L. Lions (Eds.), *Handbook of Numerical Analysis* 5 (1997) 3-207. North-Holland.