Shape Preservation During Digitization: Tight Bounds Based on the Morphing Distance

Peer Stelldinger and Ullrich Köthe

Cognitive Systems Group, University of Hamburg,
Vogt-Köln-Str. 30, D-22527 Hamburg, Germany

Abstract. We define *strong r-similarity* and the *morphing distance* to bound geometric distortions between shapes of equal topology. We then derive a necessary and sufficient condition for a set and its digitizations to be r-similar, regardless of the sampling grid. We also extend these results to certain gray scale images. Our findings are steps towards a theory of shape digitization for real optical systems.

1 Introduction

In order to make image analysis algorithms more reliable it is desirable to rigorously prove their properties, if possible. As *object shape* is of particular interest, we would like to know to what extend the information derived from a digitized shape carries over to the analog original. In this paper we analyse under which circumstances analog and digital shapes will have the same topology and bounded geometric distortions. For simplicity, we ignore the effect of gray-level quantization, and only consider spatial sampling.

The problem of topology preservation was first investigated by Pavlidis [4]. He showed that a particular class of binary analog shapes (which we will call r-regular shapes, cf. definition 6) does not change topology under discretization with any sufficiently dense square grid. Similarly, Serra [6] showed that the homotopy tree of r-regular sets is preserved under discretization with any sufficiently dense hexagonal grid. Both results apply to binary sets and the so called *subset digitization*, where a pixel is considered part of the digital shape iff its center is element of the given set. Latecki et al. [2] generalized the findings of Pavlidis to other digitizations including the *square subset* and *intersection* digitizations.

Geometric similarity can be measured by the Hausdorff distance. Pavlidis [4] proved a rather coarse bound for the Hausdorff distance between the analog and digital sets. Ronse and Tajine [5] showed that in the limit of infinitely dense sampling the Hausdorff distance between the original and digitized shapes converges to zero. However, they did not analyse topology changes.

In this paper, we combine topological and geometric criteria into two new shape similarity measures, *weak* and *strong r-similarity.* We prove that r-regularity is a necessary and sufficient condition for an analog set (i.e. a binary image) to be reconstructible (in the sense of both measures) by any regular or irregular grid with sampling distance smaller than r. These findings also apply to certain gray-scale images that result from blurring of binary images.

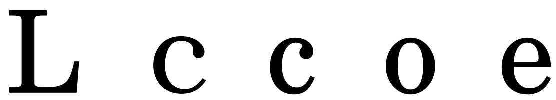

Fig. 1. Examples for weak similarity: L and c have the same topology, but large Hausdorff distance d_H. The two c's have a much smaller d_H (when overlaid). Between c and o, d_H is still quite small, but they differ by topology. The distinction between o and e is not so clear: Their topology is equal, and d_H still relatively small.

2 Shape Similarity and Digitization

Given two sets A and B, their similarity can be expressed in several ways. The most fundamental is topological equivalence. A and B are topologically equivalent if there exists a homeomorphism, i.e. a bijective function $f : A \to B$ with f and f^{-1} continuous. However, this does not completely characterize the topology of a set embedded in the plane $\mathbb{R}^2$. Therefore, [6] introduced the homotopy tree which encodes whether some components of A enclose others in a given embedding. Both notions are captured simultaneously when the homeomorphism f is extended to the entire $\mathbb{R}^2$ plane. Then it also defines a mapping $A^c \to B^c$ for the set complements and ensures preservation of both the topology and the homotopy tree. We call this an $\mathbb{R}^2$*-homeomorphism*.

Geometric similarity between two shapes can be measured by the *Hausdorff distance* d_H between the shape boundaries. We call the combination of both criteria *weak r-similarity*:

Definition 1. *Two bounded subsets* $A, B \subset \mathbb{R}^2$ *are called* weakly r-similar *if there exists a homeomorphism* $f : \mathbb{R}^2 \to \mathbb{R}^2$ *such that* $\boldsymbol{x} \in A \Leftrightarrow f(\boldsymbol{x}) \in B$, *and the Hausdorff distance between the set boundaries* $d_H(\partial A, \partial B) \leq r \in \mathbb{R}_+ \cup \{\infty\}$.

In many cases, weak r-similarity captures human perception quite nicely. Fig. 1 demonstrates this for the shape of some letters. However, it is not always sufficient. This can already be seen by comparing the letters “o” and “e” in fig. 1, but fig. 2 shows even more striking examples: These sets are topologically equivalent and have small Hausdorff distance, yet the shapes and meanings are percieved as very different. The reason for the failure of weak r-similarity is the following: Topology preservation is determined by an $\mathbb{R}^2$-homeomorphism that maps each point of A to a unique point of B. In contrast, the Hausdorff distance is calculated by mapping each point of A to the nearest point of B and vice versa. Both mappings are *independent* of each other and, in general, totally different. We can improve the similarity measure by using the same transformation for both the determination of topological equivalence and geometric similarity:

Definition 2. *Two sets* A, $B \subset \mathbb{R}^2$ *are called* strongly r-similar *and we write* $A \overset{r}{\approx} B$, *if they are weakly r-similar and* $\forall \boldsymbol{x} \in \partial A : |\boldsymbol{x} - f(\boldsymbol{x})| \leq r$. *Such a restricted homeomorphism is called* r-homeomorphism. *The* morphing distance *is defined as* $d_M(A, B) := \inf\Big(\{\infty\} \cup \{r | A \overset{r}{\approx} B\}\Big)$.

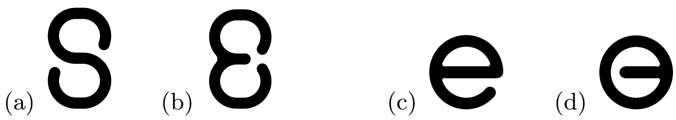

Fig. 2. Failure of weak similarity: (a) and (b) have the same topology and small Hausdorff distance, but large morphing distance. Their shapes and symbolic meanings (s vs. ε) are perceived as very different. Likewise (c) and (d).

Two sets are $\mathbb{R}^2$-topologically equivalent iff they are (strongly or weakly) ∞-similar, and they are equal iff they are (strongly or weakly) 0-similar. The morphing distance is symmetric, and the triangle inequality holds because the triangle inequality of the Euclidien metric applies at every point when two transformations are combined. Therefore, d_M is a metric and an upper bound for the Hausdorff metric d_H. It is easy to show that the existence of an r-homeomorphism that maps ∂A to ∂B implies the existence of an r-homeomorphism for the whole plane. Under strong r-similarity, the topology is not only preserved in a global, but also in a local manner: When we look at the embedding of A into $\mathbb{R}^2$ within a small open region U_A, a corresponding open region U_B with the same topological characteristics exists in the embedding of B, and the distance between the two regions is not greater than r. The shapes in fig. 2 are examples for non-local topology preservation: Morphing of corresponding shapes onto each other requires a rather big r, and $d_H \ll d_M$ in these cases.

Consider a set $A \in \mathbb{R}^2$. Its subset discretization is obtained by storing set inclusion information only at a countable number of points, i.e. on a grid:

Definition 3. *A countable set* $S \subset \mathbb{R}^2$ *of* sampling points *where* $d_H(\{\boldsymbol{x}\}, S) \leq r$ *holds for all* $\boldsymbol{x} \in \mathbb{R}^2$ *is called an* r-grid *if for each bounded set* $A \in \mathbb{R}^2$ *the subset* $S \cap A$ *is finite. The associated Voronoi regions are called* pixels*:*

$$\mathrm{Pixel}_S : S \to \mathcal{P}(\mathbb{R}^2), \quad \mathrm{Pixel}_S(\boldsymbol{s}) := \{\boldsymbol{x} : \forall \boldsymbol{s'} \in S \setminus \{\boldsymbol{s}\} : |\boldsymbol{x} - \boldsymbol{s}| \leq |\boldsymbol{x} - \boldsymbol{s'}|\}$$

The intersection of A *with* S *is called the* S-digitization of A*:* $\mathrm{Dig}_S(A) := A \cap S$.

This definition is very broad and captures even irregular grids, provided their Voronoi regions have bounded radius, see fig. 3.

As it is not useful to directly compare discrete sets with analog ones, we reconstruct an analog set from the given digitization. This is done by assigning the information stored at each sampling point to the entire surrounding pixel:

Definition 4. *Given a set* $A \subseteq \mathbb{R}^2$ *and a grid* S*, the* S-reconstruction *of* $\mathrm{Dig}_S(A)$ *is defined as* $\hat{A} = \mathrm{Rec}_S(\mathrm{Dig}_S(A)) = \bigcup_{\boldsymbol{s} \in (S \cap A)} \mathrm{Pixel}_S(\boldsymbol{s})$.

The results of a reconstruction process will be considered correct if the reconstructed set $\hat{A}$ is sufficiently similar to the original set A. Formally, we get

Definition 5. *A set* $A \subseteq \mathbb{R}^2$ *is* reconstructible *by an* r*-grid* S *if the* S*-reconstruction* $\hat{A}$ *is strongly* r*-similar to* A*, i.e.* $d_M(A, \hat{A}) \leq r$.

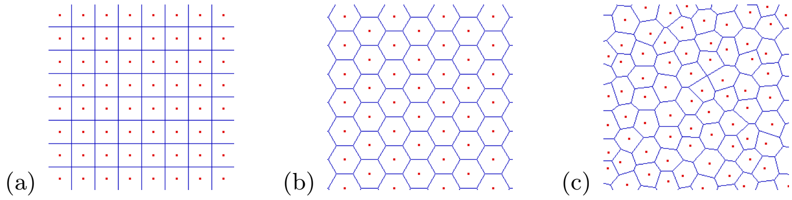

Fig. 3. Many different grid types are covered by our definition based on Voronoi regions, including regular square (a) and hexagonal (b) ones, and irregular grids (c) as found in natural image acquisition devices like the human eye.

3 Conditions for Shape Preserving Digitization

It turns out that shape preserving digitization is only possible if the shape fulfills a regularity requirement developed independently by Pavlidis [4] and Serra [6]:

Definition 6. *A compact set $A \subset \mathbb{R}^2$ is called* r-regular *iff for each boundary point of A it is possible to find two osculating open balls of radius r, one lying entirely in A and the other lying entirely in A^c.*

Using this definition, we prove the following geometric sampling theorem:

Theorem 1. *Let $r \in \mathbb{R}_+$ and A an r-regular set. Then A is reconstructible with* any *r'-grid S, $0 < r' < r$.*

In a previous paper [1], we proved the same theorem, but reconstructible sets were defined by means of *weak* r-similarity (simply called r-similarity there). Here we will show that the theorem also holds when *strong r'-similarity* is required. Moreover, theorem 2 shows that r-regularity is also a necessary condition.

Proof. From the weak version of the theorem in [1] we already know that the reconstruction is $\mathbb{R}^2$-topologically equivalent to A, and the Hausdorff distance between the boundaries is at most r'. To tighten the theorem for strong r'-similarity it remains to be shown that there even exists an r'-homeomorphism.

Due to the r-regularity of A, no pixel can touch two components of ∂A. Therefore, we can treat each component $\partial A'$ of ∂A and its corresponding component $\partial \hat{A}'$ separately. The proof principle is to split $\partial A'$ and $\partial \hat{A}'$ into sequences of segments $\{C_i\}$ and $\{\hat{C}_i\}$, and show that, for all i, $\hat{C}_i$ can be mapped onto C_i with an r'-homeomorphism. The order of the segments in the sequences is determined by the orientation of the plane, and corresponding segments must have the same index. Then the existence of an r'-homeomorphism between each pair of segments implies the existence of the r'-homeomorphism for the entire boundary. We define initial split points $\hat{\boldsymbol{c}}_i$ of $\partial \hat{A}'$ as follows (see fig. 4a):

Case 1: A split point is defined where $\partial \hat{A}'$ crosses or touches $\partial A'$. *Case 1a:* If this is a single point, it automatically defines a corresponding split point of $\partial A'$. *Case 1b:* If extended parts of the boundaries coincide, the first and last common points are chosen as split points.

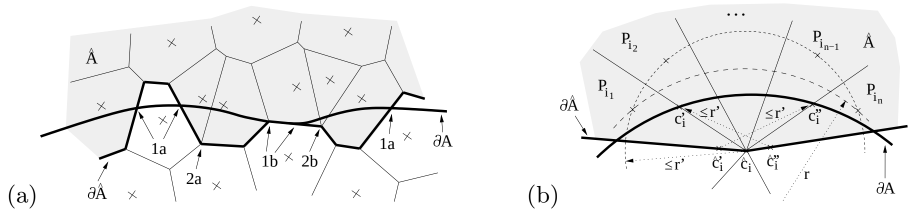

Fig. 4. (a) Different cases for the definition of split points; (b) Partition refinement and mapping for case 2b.

Case 2: A pixel corner which is on $\partial\hat{A}'$ but not on $\partial A'$ becomes a split point if the corner point lies in A (A^c) and belongs to at least two pixels that are in $\hat{A}^c$ ($\hat{A}$). *Case 2a:* If there are exactly two such neighboring pixels, a corresponding split point is defined where $\partial A'$ crosses the common boundary of these pixels. *Case 2b:* Otherwise, the split point is treated specially.

In the course of the proof, the initial partition will be refined. The treatment of case 1b is straigthforward: Here, two segments C_i and $\hat{C}_i$ coincide, so we can define the r'-homeomorphism as the identity mapping.

Next, consider case 2b (fig. 4b). Let the special split point $\hat{\boldsymbol{c}}_i \in A$ (A^c) be a corner of pixels $P_{i_1}, ..., P_{i_n} \in \hat{A}^c$ ($\hat{A}$). The orientation of the plane induces an order of these pixels. The pixels P_{i_2} to $P_{i_{n-1}}$ intersect $\partial\hat{A}'$ only at the single point $\hat{\boldsymbol{c}}_i$. We must avoid that an extended part of $\partial A'$ gets mapped onto the single point $\hat{\boldsymbol{c}}_i$. Thus, we change the initial partitioning: Replace $\hat{\boldsymbol{c}}_i$ with two new split points $\hat{\boldsymbol{c}}_i{}'$ and $\hat{\boldsymbol{c}}_i{}''$, lying on $\partial\hat{A}'$ to either side of $\hat{\boldsymbol{c}}_i$ at a distance ε. Define as their corresponding split points the points $\boldsymbol{c}_i'$ and $\boldsymbol{c}_i''$ where $\partial A'$ crosses the commen border of P_{i_1}, P_{i_2} and $P_{i_{n-1}}, P_{i_n}$ respectively. Due to r-regularity, $|\overline{\boldsymbol{c}_i'\hat{\boldsymbol{c}}_i}| < r'$ and $|\overline{\boldsymbol{c}_i''\hat{\boldsymbol{c}}_i}| < r'$, and the same is true for all points between $\boldsymbol{c}_i'$ and $\boldsymbol{c}_i''$. Therefore, ε can always be chosen so that every point between $\boldsymbol{c}_i'$ and $\boldsymbol{c}_i''$ can be mapped onto every point between $\hat{\boldsymbol{c}}_i{}'$ and $\hat{\boldsymbol{c}}_i{}''$ with a displacement of at most r'. This implies the existence of an r'-homeomorphism between these segments.

After these modifications, the segments not yet treated have the following important properties: Each C_i is enclosed within one pixel P_i, and the corresponding segment $\hat{C}_i$ is a subset of P_i's boundary. To prove the theorem for these pairs, we use the property of Reuleaux triangles with diameter r' that no two points in such a triangle are farther apart than r' (fig. 5a). Due to r-regularity, $\partial A'$ can cross the border of any r'-Reuleaux triangle at most two times. We refine the segments so that each pair is contained in a single triangle, which implies the existence of an r'-homeomorphism. Consider the pair C_i, $\hat{C}_i$ and let the sampling point of pixel P_i be $\boldsymbol{s}_i$. If this point is not on $\partial A'$ (fig. 5b), C_i splits P_i into two parts, one containing $\hat{C}_i$ and the other containing $\boldsymbol{s}_i$. We now place r'-Reuleaux triangles as follows: a corner of every triangle is located at $\boldsymbol{s}_i$, every triangle intersects C_i and $\hat{C}_i$, and neighboring triangles are oriented at 60° of each other, so that no three triangles have a common overlap region. Since the pixel radius is at most r', this set of triangles completely covers both

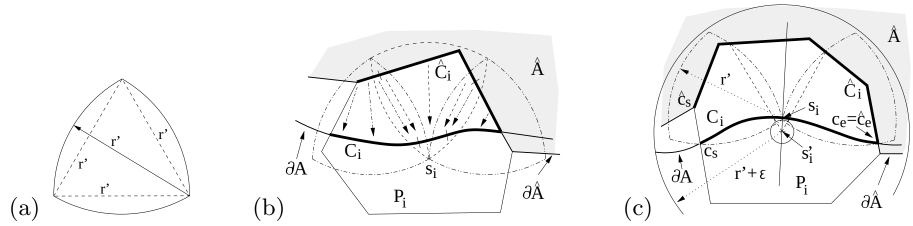

Fig. 5. (a) Any two points in a Reuleaux triangle of size r' have a distance of at most r'; (b) Covering of corresponding segments with Reuleaux triangles; (c) Construction for sampling points lying on $\partial A'$.

C_i and $\hat{C}_i$, and each consecutive pair of triangles shares at least one point of either segment. Thus, we can define additional split points among the shared points, so that corresponding pairs of the new segments lie entirely within one triangle. The existence of an r'-homeomorphism for the refined segments follows.

If the sampling point $\boldsymbol{s}_i$ of P_i is on $\partial A'$ (fig. 5c), this Reuleaux construction does not generally work. In this case, we first place two r'-Reuleaux triangles such that both have $\boldsymbol{s}_i$ as a corner point, one contains the start points $\boldsymbol{c}_s, \hat{\boldsymbol{c}}_{\boldsymbol{s}}$ of C_i and $\hat{C}_i$ respectively, the other the end points $\boldsymbol{c}_e, \hat{\boldsymbol{c}}_{\boldsymbol{e}}$, and they overlap $\hat{C}_i$ as much as possible. If they cover $\hat{C}_i$ completely, the Reuleaux construction still works with $\boldsymbol{s}_i$ as split point. Otherwise $\hat{C}_i$ is partly outside of the triangles, and the normal of $\partial A'$ crosses $\hat{C}_i$ in this outside region. We choose a point $\boldsymbol{s}'_i$ on the opposite normal with distance ε from $\boldsymbol{s}_i$ and project each point $\boldsymbol{c}$ of $\hat{C}_i$ not covered by either triangle onto the point where the line $\overline{\boldsymbol{c}\boldsymbol{s}'_i}$ crosses C_i. It can be seen that this mapping is an r'-homeomorphism: Draw circles with radius ε and $r' + \varepsilon$ around $\boldsymbol{s}'_i$. C_i and $\hat{C}_i$ lie between these circles, so that each point is moved by at most r'. The extreme points of this construction define new split points, and the remaining parts of C_i and $\hat{C}_i$ can be mapped within the two triangles. Thus, there is an r'-homeomorphism in this case as well. □

As the proof makes no assumptions about pixel shape, the geometric sampling theorem applies to *any* regular or irregular r'-grid (cf. fig 3). Moreover, when a set is resonstructible by some grid S, this automatically holds for any translated and rotated copy of S as well. r-regularity is not only a sufficient but also a nessessary condition for a set to be reconstructible:

Theorem 2. *Let A be a set that is not r-regular. Then there exists an r'-grid S with $0 < r' < r$ such that A is not reconstructible by S.*

Proof. We explicitly construct such a grid. There are two cases: *Case 1*: If A is not r''-regular for any $r'' > 0$, then it contains at least one corner or junction. In both cases it is possible to place sampling points so that the reconstruction of a connected set becomes disconnected, and the topology is not preserved (fig. 6 a and b). *Case 2*: Let A be r''-regular with $0 < r'' < r' < r$. Then there is a maximal inside or outside circle of radius r'' with center $\boldsymbol{p}_0$ that touches ∂A in

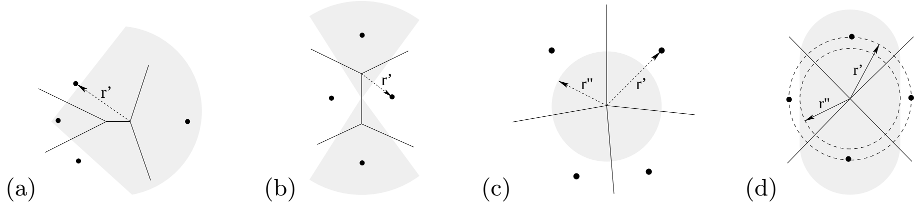

Fig. 6. Examples where the topology of the reconstruction by an r-grid (balls and lines) differs from the topology of the original set (gray area) because it is not r-regular.

at least two points. Draw a circle with radius r' around $\boldsymbol{p}_0$. *Case 2a*: If the r''-circle coincides with a component of ∂A, a component of A or A^c is completely inside the r'-circle, and we can place sampling points on this circle such that the enclosed component is lost in the reconstruction. *Case 2b*: Otherwise, an r' can be chosen so that part of the r'-circle is in A, part in A^c. If these parts form more than two connected components, we can place a sampling point in each component, and the reconstruction contains a junction whereas the original shape does not. *Case 2c*: If there are exactly two components, we can move the r'-circle a little so that it will either no longer intersect with ∂A, which brings us back to case 2a, or the number of components will increase, which brings us to case 2b. In any case, the topology of A is not preserved (fig. 6 c and d). □

The geometric sampling theorems do not only hold for binary images, but also for all r-regular level sets of gray-level images. In particular, we proved in [1] that theorem 1 also applies if r-regular sets are subjected to blurring with a circular averaging filter before digitization. Such images are approximations of what could be observed with a real camera (albeit real point spread functions are more complicated). We proved the following theorem:

Theorem 3. *Let A be an r-regular set, k_p a circular averaging filter with radius $p < r$, and $f_A = k_p \star \chi_A$ the blurred image of A (χ_A is A's characteristic function, $\star$ denotes convolution). Further let L_l be any level set of f_A and S an r''-grid with $r'' < r - p$. Then the S-reconstruction $\hat{L}_l$ of L_l is $(p + r'')$-similar to A.*

In [1], reconstruction referred to weak r-similarity, but the theorem can be extended to strong r-similarity. The original proof first showed that any level set L_l of f_A is $(r - p)$-regular and p-similar to A. Then it followed from theorem 1 that L_l is reconstructible by any grid with pixel radius $r'' < r - p$. Since theorem 1 has been tightened for strong r-similarity, it remains to be shown that the first part of the proof can also be tightened:

Proof. We must show that there exists a p-homeomorphism between A and any level set L_l after blurring. We already know that there is an $\mathbb{R}^2$-homeomorphism. Consider a p-wide strip A_p around ∂A. Due to r-regularity, the normals of ∂A cannot cross within A_p. Therefore, every point in A_p is crossed by exactly one normal, and the starting point of the normal is at most at distance p. Since the level lines ∂L_l always run in the strip A_p without crossing any normal twice (see [1], lemma 5), we can define a p-homeomorphism from ∂A to any level line by mapping each point along its normal. □

4 Findings

In this paper we proved a powerful geometric sampling theorem. In intuitive terms, our theorem means the following: When an r-regular set is digitized, its boundaries move by at most half the pixel diameter. The number of connected components of the set and its complement are preserved, and the digital sets are directly connected (more precisely, they are well-composed in the sense of Latecki [3]). Parts that were originally connected do not get separated. As these claims hold for any regular or irregular r'-grid, they also hold for translated and rotated versions of some given grid. Thus, reconstruction is robust under Euclidian transformations of the grid or the shape.

Since strong r-similarity is also a local property, we can still apply our results if a set is not r-regular in its entirety. We call a segment of a set's boundary *locally r-regular* if definition 6 holds at every point of the segment. Theorem 1 then applies analogously to this part of the set because the boundary segment could be completed into some r-regular set where the theorem holds everywhere, and in particular in a local neighborhood of the segment.

Our results can be generalized to gray-level images in two ways: First, they apply to all level sets or parts thereof that are r-regular. This is usually the case at edge points that are sufficiently far from other edges or junctions. Second, when a binary image is first blurred by a circular averaging filter, the theorem still holds with $r = r' + p$, where r' and p are the radii of the pixels and filter respectively. This is similar to Lateckis work ([2,3]), as his v-digitization amounts to blurring with a square averaging filter. Our findigs are important steps towards a geometric sampling theory applicable to real optical systems. In the future, we will try to extend them to more general filter classes. We will also analyse what happens in the neighborhood of junctions (where the r-regularity constraint cannot hold), and under the influence of gray-level quantization and noise.

References

1. Köthe U., Stelldinger, P.: *Shape Preserving Digitization of Ideal and Blurred Binary Shapes*, Univ. Hamburg, Informatics Dept., Tech. Rep. FBI-HH-M-325/03, 2003.
2. Latecki, L.J., Conrad, C., Gross, A.: *Preserving Topology by a Digitization Process.* Journal of Mathematical Imaging and Vision **8**, 131–159, 1998.
3. Latecki, L.J.: *Discrete Representation of Spatial Objects in Computer Vision.* Kluwer Academic Publishers, Dordrecht, 1998.
4. Pavlidis, T.: *Algorithms for Graphics and Image Processing.* Computer Science Press, Rockville, Maryland. 1982.
5. Ronse, C., Tajine, M.: *Discretization in Hausdorff Space.* Journal of Mathematical Imaging and Vision **12**, 219–242, 2000.
6. Serra, J.: *Image Analysis and Mathematical Morphology* Academic Press, New York, 1982.