Towards a General Sampling Theory for Shape Preservation

Peer Stelldinger [a] Ullrich Köthe [a]

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

Abstract

Computerized image analysis makes statements about the continous world by looking at a discrete representation. Therefore, it is important to know precisely which information is preserved during digitization. We analyse this question in the context of shape recognition. Existing results in this area are based on very restricted models and thus not applicable to real imaging situations. We present generalizations in several directions: first, we introduce a new shape similarity measure that approximates human perception better. Second, we prove a geometric sampling theorem for arbitrary dimensional spaces. Third, we extend our sampling theorem to 2-dimensional images that are subjected to blurring by a disk point spread function. Our findings are steps towards a general sampling theory for shapes that shall ultimately describe the behavior of real optical systems.

Key words: shape preservation, digitization, discretization, sampling theory, topology, Hausdorff distance, r-regularity

1 Introduction

Computerized image analysis is increasingly used in applications where errors and mistakes can have critical consequences, like industrial inspection, surveillance, autonomous vehicle control, and medical imaging. Therefore, it becomes more and more important to understand and formalize the conditions of successful algorithm behavior. Errors may have many causes, some of them being external like viewing conditions and illumination, but others depending on the system itself. Since one has often no control over the external influences, it is all the more important to study how certain internal design choices influence performance limits of image analysis systems.

One key aspect is the role of sampling: What can be seen with a digital device looking at a continous world? In this paper we will be concerned with the relationship of sampling and shape recognition: How should an image analysis process be designd in order to preserve important shape characteristics of real objects, and when can this performance be guaranteed? Ultimately, this question should be answered by a general geometric sampling theorem analogous to Shannon's famous sampling theorem describing the preservation of waveforms.

Naturally, this paper is not the first attempt to answer this question, neither will it be the last. A detailed description of previous work is given in the next section. This work is characterized by the fact that definite results could only be obtained by imposing very restrictive assumptions on the digitization process. The results are therefore quite far from reality and provide relatively little guidance for the design of practical applications. It is our goal to extend the existing work in ways that eventually bridge the gap between theory and praxis. In particular, we are going to present an improved measure of shape similarity (section 3), we present the class of sets we are working on (section 4), we define which digitization results have to be seen as correct under a certain digitization model (section 5), we generalize existing sampling theorems for binary shapes to multiple dimensions and arbitrary grids (section 6), we prove tight bounds for the 2-dimensional case (section 7), and we analyse how the sampling theorem has to be modified when the image is subjected to blurring prior to digitization, since blurring is unavoidable in any real optical system due to diffraction effects, defocus and finite sensor size (section 8). These extensions represent significant steps towards a truly general geometric sampling theorem, although further generalizations are still needed to achieve the ultimate goal.

2 Prior Work

Since image analysis tries to derive knowledge about the real world by analysing discrete representations, it is important to understand which information is preserved during digitization. In particular we are interested in the preservation of shape characteristics. There exist several approaches which describe how shapes – which can be understood as binary sets – behave under certain types of digitization. First results on this problem were independently presented by Pavlidis [5] and Serra [7] in 1982. Both discovered the fundamental role of r-regular sets, which we are reviewing in detail in section 4.

Pavlidis [5] showed that 2-dimensional r-regular shapes (cf. definition 3) do not change their topology under discretization with any square grid having a minimal sampling point distance less than $\sqrt{2}r$. In this context digitization means that a pixel is part of the digital set iff its center point lies in the original continuous shape (this digitization is often called subset digitization, see [4]). Similarly, Serra [7] proved, that the homotopy tree of 2D r-regular shapes does not change under subset digitization with any hexagonal grid where the minimal sampling point distance is at most r.

Both Pavlidis and Serra only considered the 2D case. In contrast, Ronse and Tajine [6] showed that in any dimension the Hausdorff discretization of a set with any sampling grid type converges to the original set if the sampling distance converges to zero. In addition, they showed for the 2D case and square grids that the topology of r-regular shapes is also preserved under Hausdorff discretization if the minimal sampling point distance is at most $r/2$, see [10]. Unfortunately Hausdorff discretizations are very inconvenient for practical applications because they are defined globally and heavily depend on the alignment of the grid relatively to the set.

Latecki et al. [4,3] extended Pavlidis' results by replacing subset digitization with a more general but still locally defined concept called v-digitization. Here a pixel is part of the digital set, iff the percentage of the pixel area covered by the original shape is at least $v \in [0\% \ldots 100\%]$. They proved that the topology of a 2-dimensional r-regular image is preserved if it is sampled with a square grid with minimal sampling point distance less than $r/\sqrt{2}$.

In 1999, Giraldo et al. [1] proved that any compact set, which has a finite polyhedron as boundary can be digitized by a square grid of a certain density, such that the digitization is homotopically equivalent to the original one after applying a simple, well defined repairing method. However, this is only possible for polyhedra and they did not say anything about topology preservation.

By comparing the above results one recognizes that apart from the last one they all arrive at somewhat different conclusions starting from the same concept of r-regular shapes. This suggests that r-regularity is a fundamental property whose power has not been fully exploited by the existing approaches. On the other hand the requirements imposed by the existing theorems are very special and can never be met by any real imaging device. This refers mainly to the use of binary sets which are not observable in reality due to the inherent blurring of real optical systems.

In two previous papers we presented ideas to overcome these problems. In [2] we combined previously used shape similarity criteria (preservation of topology and homotopy trees, bounded Hausdorff distance) and proved that the sampling theorems of Pavlidis and Serra still hold under these stricter conditions, even if arbitrary grids are used. We also showed how the sampling density has to be changed for a particular class of blurred shapes. In [9] we introduced an even stricter similarity criterion and proved that r-regularity is not only necessary but also sufficient for correct digitization of binary shapes under the new criterion. In the present paper we integrate these results for the first time. In addition, we extend our definitions to arbitrary dimensions and analyze the possiblility to prove an n-dimensional sampling theorem for non-blurred shapes.

3 Shape Similarity

Due to the information loss during digitization, a digital shape is almost never identical to its continuous original. In order to formally describe this information loss, a precise definition of shape similarity is needed. Ideally this definition should resemble human perception as good as possible. As we will see, previously used similarity criteria do not always fulfill this requirement. Therefore we will define a new criterion called *strong r-similarity* that will be the basis of our proofs. Since shapes can be considered as closed subsets of $\mathbb{R}^2$, we need a criterion to compare 2-dimensional sets. Pavlidis [5] suggested that the sets should be topologically equivalent. Two sets A and B are topologically equivalent if there exists a homeomorphism $f : A \to B$, i.e. a bijective function with f and f^{-1} continuous. Unfortunately, topological equivalence imposes no restrictions on the space where this homeomorphism is defined. This means that 2D objects can be topologically equivalent even if such a homeomorphism exists only for a 3-dimensional embedding. Consequently two 2D sets can be topologically equivalent whithout having the same inclusion properties for their components in $\mathbb{R}^2$, as can be seen in fig. 1(a). Of course this problem may arise for shapes of any dimension. It can be solved by requiring the homeomorphism to be defined in a space whose dimension equals the dimension of the shapes. An alternative way to capture the inclusion properties correctly was proposed by Serra [7], who defined homotopy trees for shape comparison. A homotopy tree encodes which components of A enclose others in a given embedding into $\mathbb{R}^2$. Unfortunately identical homotopy trees do not always imply high shape similarity, as fig. 1(b) shows. While topology preservation and identity of homotopy trees is necessary, geometric shape similarity cannot be neglected. As suggested by Ronse and Tajine [6], geometric shape similarity can be measured by the *Hausdorff distance*

$$d_H(\partial A, \partial B) = \max \left(\max_{\vec{x} \in \partial A} \min_{\vec{y} \in \partial B} d(\vec{x}, \vec{y}), \max_{\vec{y} \in \partial B} \min_{\vec{x} \in \partial A} d(\vec{x}, \vec{y}) \right)$$

between the shapes' boundaries ∂A and ∂B of A and B, where $d(\cdot, \cdot)$ is a metric of the space (in the sequel we will always use the Euclidean metric). But a small Hausdorff distance is also not a sufficient similarity criterium, as fig. 1(c) shows, where an S has changed into a number 6 without the two characters having large Hausdorff distance.

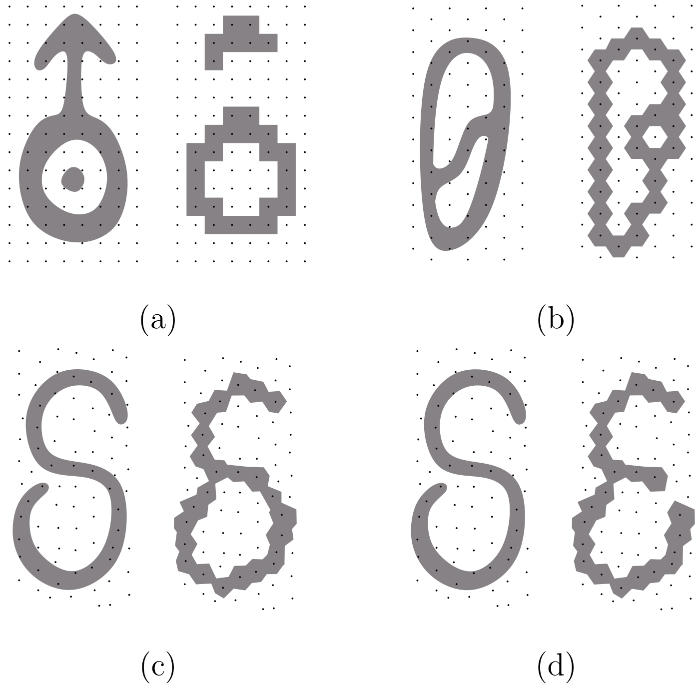

Figure 1. Comparison of similarity criteria. (a): The shapes are homeomorphic and thus topologically equivalent. But the homeomorphism can not be embedded in $\mathbb{R}^2$. (b): the shapes have the same homotopy tree (the infinite background component is connected to one foreground component, which is connected to two finite background components), but are not percieved as similar. (c): The Hausdorff distance of the two shapes is less than the pixel radius of the grid, but the shapes are topologically different. (d): The shapes are topologically equivalent in $\mathbb{R}^2$ and have small Hausdorff distance, but are still perceived as very different.

Since topological equivalence would have captured this error it is logical to combine both criteria. We call this *weak r-similarity*. For generality we define this concept for arbitrary dimensions as follows:

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

Note that we require the homeomorphism to exist in the entire space $\mathbb{R}^n$, which means that not only the shapes but also their complements have to be topologically equivalent. This implies identity of homotopy trees whereas the converse is not true for $n \geq 3$, as the example of two tori passing through each other vs. two separated tori illustrates.

Intuitively the homeomorphism can be understood as a transformation of an elastic, rubber-like space, which imposes a smooth distortion such that A gets equal to B. We call this an $\mathbb{R}^n$*-homeomorphism.*

L c c o e

Figure 2. 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.

In many cases, weak r-similarity captures human perception quite nicely. Fig. 2 demonstrates this for the shape of some letters in 2D. Other examples are 1(a) and (b), where the shapes are only weakly r-similar for relatively great r. However, it is not always sufficient. This can already be seen by comparing the letters “o” and “e” in fig. 2, but fig. 1(d) shows an even more striking example: These sets are weakly r-similar for very small r, so they are topologically equivalent and have the same homotopy trees and a small Hausdorff distance. Nevertheless the shapes and meanings are percieved as very different. The failure of weak r-similarity is because of our definition of the geometric distance: While topology preservation is determined by an $\mathbb{R}^n$-homeomorphism that maps each point of A to a unique point of B, the Hausdorff distance is calculated by a completely different mapping of each point of A to the nearest point of B and vice versa.

In order to improve the similarity measure, we require the same transformation to be used for both the determination of topological equivalence and geometric similarity. We call this *strong r-similarity*:

Definition 2 *Two sets A, $B \subset \mathbb{R}^n$ are called* strongly r-similar *and we write $A \overset{r}{\approx} B$, if there exists a homeomorphism $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ such that $\vec{x} \in A \Leftrightarrow f(\vec{x}) \in B$ and $\forall \vec{x} \in \partial A : |\vec{x} - f(\vec{x})| \leq r$. Such a restricted homeomorphism will be called* r-homeomorphism *in $\mathbb{R}^n$ and the sets will be called r-topologically equivalent. The* morphing distance *is defined as $d_M(A, B) := \inf\Big(\{\infty\} \cup \{r | A \overset{r}{\approx} B\}\Big)$.*

Since the concatenation of an r-homeomorphism with an s-homeomorphism is obviously at least an $(r+s)$-homeomorphism, the morphing distance is a metric (see also [9]). Since d_M is an upper bound for the Hausdorff distance, strong r-similarity implies weak r-similarity. Under strong r-similarity, the topology is not only preserved in a global, but also in a local manner: When looking at the embedding of A into $\mathbb{R}^n$ 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 (a),(b) and (d) in fig. 1 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, thus they are only strongly r-similar for relatively

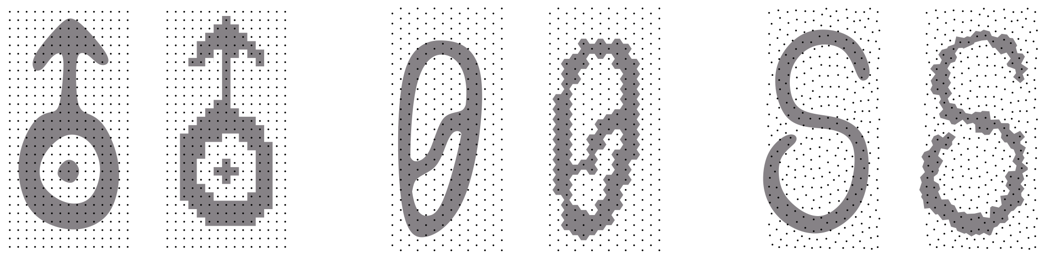

(a) (b) (c)

Figure 3. Using sufficiently dense sampling grids. When sampling the shapes (a), (b) and (c) with grids which are dense enough, their reconstructions are similar in the sense of human perception and they are also strongly r-similar.

great r. By using a denser sampling grid, the similarity can be improved, such that r is not greater than the pixel radius. If this is done, shape characteristics are really preserved in the sense of human perception, as can be seen in fig. 3.

Note that the definition of r-homeomorphisms is similar, but not equal to the definition of ε-homeomorphisms used in the work of Giraldo et al. [1]. The ε-homeomorphism is only defined on the two compared sets and not in the background as well. Thus the important inclusion properties of components are not captured.

4 r-regular Sets

A shape can be defined as a set $A \in \mathbb{R}^n$. A geometric sampling theorem states which type of shapes retains characteristic properties during digitization. Surprisingly, in 1982 Pavlidis [5] and Serra [7] independently introduced totally different definitions for basically the same type of sets, which they used in their theorems. Also Latecki et al. [3,4] and Tajine and Ronse [10] referred to this shape class. Since this shape class is fundamental, we introduce here a unified definition of *r-regular sets* (see fig. 4), which can be proved to be equivalent to the other definitions for the 2D case (see [8]).

Definition 3 *A compact set $A \subset \mathbb{R}^n$ 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.*

If a set is r-regular, each part of its shape has a minimal width (diameter) of $2r$. Moreover, for each boundary point there exists one unique tangent hyperplane, and the curvature is defined for every tangent direction and bounded by $\frac{1}{r}$. r-regular sets can be characterized as exactly those compact sets, that are morphologically open and closed with respect to an open r-ball structuring element.

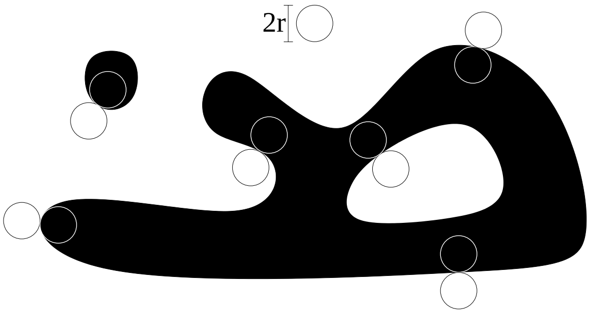

Figure 4. A 2D r-regular set.

5 Sampling and Reconstruction

In this section we develop a precise definition of the digitization process. All definitions are given for arbitrary dimensional spaces, so that subsequent theorems can be formulated as general as possible. Since we identify shapes with subsets $A \in \mathbb{R}^n$ and sets can be interpreted as binary functions, a set $A \subseteq \mathbb{R}^n$ can be transformed into an *analog binary image* by means of its characteristic function $\chi_A : \mathbb{R}^n \to \{0, 1\}$, $\chi_A(\vec{x}) = 1$ iff $\vec{x} \in A$. Now the discretization is obtained by storing the values of this image only at a countable number of sampling points. Formally we get:

Definition 4 *A countable set $S \subset \mathbb{R}^n$ of points with $d_H(\mathbb{R}^n, S) \leq r$ for some $r \in \mathbb{R}_+$ such that for each bounded set A the subset $S \cap A$ is finite, is called* r-grid*. The elements of S are the* sampling points*, and the* pixels *are defined as their associated Euclidean Voronoi regions:*

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

*Thus pixels are topological n-cells and bounded by $(n-1)$-cells, which in turn are bounded by $(n-2)$-cells and so on, eg. in 2D they are 2-cells, their edges are 1-cells and their corners are 0-cells. The intersection of $A \subseteq \mathbb{R}^n$ with S is called the S-*digitization *of A, and the restriction of the domain of A's characteristic function to S is the associated* digital binary image*:*

$$\text{Dig}_S(A) := A \cap S$$
$$\text{DigitalImage}_S(\chi_A) := \chi_A|_S : S \to \{0, 1\}$$

This definition captures arbitrary grids of any dimension. So not only square, hexagonal, cubic or other regular grids, but also irregular grids are permitted, provided the pixel shapes can be defined as Voronoi regions with bounded radius, see fig.5. For simplicity we will use the term *pixel* also in dimensions greater than two. Since it makes no sense to compare the topological properties of an analog shape (region) with a discrete set of points we first have to reconstruct an analog shape form the digital representation. We do this by assigning

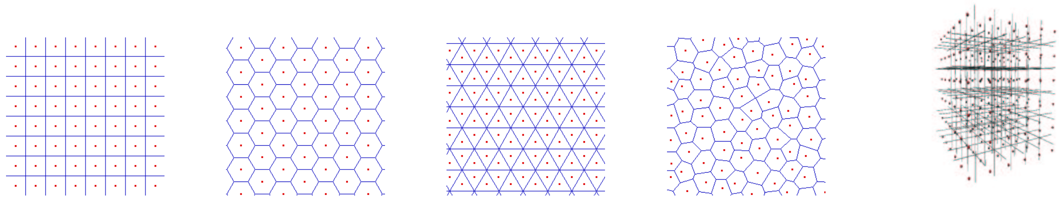

(a) (b) (c) (d) (e)

Figure 5. Many different grid types can be described when pixels are defined as the Voronoi regions of suitably located sampling points. In 2D, these include regular grids like the square (a), hexagonal (b) and trigonal ones (c), and irregular grids (d) as found in natural image acquisition devices like the human eye. The definition also captures grids of any dimension, e.g. cubic grids in 3D (e).

the information stored at each sampling point to the entire surrounding pixel:

Definition 5 *Given a set $A \subseteq \mathbb{R}^n$ and a grid S, the S-*reconstruction *of* $\mathrm{Dig}_S(A)$ *is defined as*

$$\hat{A} = \mathrm{Rec}_S(\mathrm{Dig}_S(A)) = \bigcup_{\vec{s} \in (S \cap A)} \mathrm{Pixel}_S(\vec{s})$$

Now we consider a reconstruction as correct if it is strongly r'-similar to the original shape for some small r'. Since the maximal pixel radius of an r-grid S is r, it is natural to require $r' = r$. This is the strongest useful bound since if the sampling point of one pixel with maximum radius lies on the shape's boundary, the Hausdorff distance between the shape and its reconstruction can obviously equal r. So the highest degree of similarity that can be achieved regardless of the alignment of the r-grid is strong r-similarity.

Definition 6 *A set $A \subseteq \mathbb{R}^n$ is* reconstructible *by an r-grid S if the S-reconstruction $\hat{A}$ is strongly r-similar to A, i.e. $d_M(A, \hat{A}) \leq r$. A is called* weakly reconstructible *by S if $\hat{A}$ is weakly r-similar to A.*

6 Sampling in Arbitrary Dimensions

In the following we will show that in any dimension r-regular sets are an important class for shape preserving digitization. In 2D we will show that an r-regular set is weakly reconstructible by any r'-grid with $r' < r$, regardless of the grid structure and alignment. Nevertheless, it is not possible to prove a reconstruction to be r-similar to an original set of higher dimension in general. Thus, digitization in higher dimensions is not as easy as it looks like in 2D.

The following lemmas describe some prerequisites. We only formulate them for the foreground A, but their claims and proofs apply to the background A^c analogously.

Lemma 7 *Let A be an r-regular set and $\hat{A}$ the reconstruction of A by an r'-grid S, with $0 < r' < r$. Then two sampling points lying in different components of A cannot lie in the same component of $\hat{A}$.*

PROOF. The Hausdorff distance of two components of an r-regular set A is at least $2r$, because each boundary point has an osculating r-ball lying entirely outside each component. Since the S-reconstruction of any component A' is a subset of the r'-dilation of A', the Hausdorff distance between two components of $\hat{A}$ is at least $2r - 2r' > 0$. Thus the reconstruction process cannot merge two components of A. $\square$

Lemma 8 *Let A' be a component of an r-regular set A, S be an r'-grid, $0 < r' < r'' < r$. Further, let $A'_\ominus = (A' \ominus \overline{\mathcal{B}}_{r''})^0$ be the open interior of the erosion of A' with a closed ball of radius r'', and $S_i := \{\vec{s} \in S : \mathrm{Pixel}(\vec{s}) \cap A'_\ominus \neq \emptyset\}$ the set of all sampling points whose pixels intersect $A'_\ominus$. Then at least one sampling point from S_i is in A'.*

PROOF. Since A is r-regular, every component A' contains at least one ball of radius r. The center $\vec{m}$ of such a ball lies in $A'_\ominus$. Let $\vec{s} \in S$ be a sampling point with $\vec{m} \in \mathrm{Pixel}(\vec{s})$. Then $\vec{s}$ is also element of S_i and the distance between $\vec{s}$ and $\vec{m}$ is at most $r' < r''$. Thus, $\vec{s}$ lies within A'. $\square$

Lemma 9 *Let A, A', S and S_i be defined as in lemma 8. Then any pair of pixels with sampling points in S_i is connected by a chain of adjacent pixels whose sampling points are also in S_i. Pixels are adjacent if they are direct neighbors, that is their common boundary has dimension n-1 (a hyper-face, ie. an edge in 2D).*

PROOF. Every component A' of an r-regular set A is r-regular, too. Thus $A'_\ominus$ is an open, connected set. Now let $\vec{s}_1$ and $\vec{s}_2$ be sampling points in S_i. The interior of their pixels intersects $A'_\ominus$, and there exist two points $\vec{s}_1{}'$, $\vec{s}_2{}'$ lying in $(\mathrm{Pixel}(\vec{s}_1))^0 \cap A'_\ominus$ and $(\mathrm{Pixel}(\vec{s}_2))^0 \cap A'_\ominus$ respectively. $\vec{s}_1{}'$ and $\vec{s}_2{}'$ can be connected by a path in $A'_\ominus$ which, without loss of generality, only intersects the pixel boundaries in the interior of their bounding $(n-1)$-cells. The sampling points of all pixels intersecting this path are in S_i as well. The order in which the path enters those pixels defines a chain of adjacent pixels. $\square$

Lemma 10 *Let A, A', S and S_i be defined as in lemma 8. Then each sampling point lying in A' is either a member of S_i or is connected to a member of S_i by a chain of adjacent pixels whose sampling points all lie in A'.*

PROOF. Let $\vec{c}$ be any sampling point in A'. Then there exists a ball of radius r in A' such that $\vec{c}$ lies in the ball. Let $\vec{m} \in A'_\ominus$ be the center of the ball. The halfline starting at $\vec{c}$ and going through $\vec{m}$ crosses the boundary of the convex $\text{Pixel}(\vec{c})$ at exactly one point $\vec{c}'$. If $d(\vec{c}, \vec{m}) \leq d(\vec{c}, \vec{c}')$, the point $\vec{m}$ is part of $\text{Pixel}(\vec{c})$ and thus $\vec{c} \in S_i$. If $d(\vec{c}, \vec{m}) > d(\vec{c}, \vec{c}')$, let g be the $(n-1)$-dimensional hyper-plane defined by the $(n-1)$-cell of $\text{Pixel}(\vec{c})$ whose closure contains $\vec{c}'$. If there are several such hyper-planes, one is chosen arbitrarily. Due to the definition of Voronoi regions the point $\vec{c}''$ constructed by mirroring $\vec{c}$ on g is a sampling point in S, and $\text{Pixel}(\vec{c}'')$ is adjacent to $\text{Pixel}(\vec{c})$. Since $c := d(\vec{c}', \vec{c}) = d(\vec{c}', \vec{c}'')$, the point $\vec{c}''$ always lies on the hyper-sphere of radius c with center $\vec{c}'$. Since $\vec{c}, \vec{c}'$ and $\vec{m}$ are collinear, $\vec{c}$ has the largest distance to $\vec{m}$ among all points on this hyper-sphere, and so $d(\vec{m}, \vec{c}'') < d(\vec{m}, \vec{c})$. Thus, the sampling point $\vec{c}''$ lies in A', and is closer to $\vec{m}$ than $\vec{c}$. We can repeat this construction iteratively to obtain a sequence of adjacent pixels whose sampling points successively get closer to $\vec{m}$. Since there are only finitely many sampling points in A', one such pixel will eventually intersect $A'_\ominus$. $\square$

Theorem 11 (sampling theorem for n-dimensional binary images) *Let $r \in \mathbb{R}_+$ and $A \subset \mathbb{R}^n$ an r-regular set. Further let $\hat{A}$ be the reconstruction of A with an arbitrary r'-grid S in $\mathbb{R}^n$, $0 < r' < r$. Then A and $\hat{A}$ have identical homotopy trees and the Hausdorff distance between their boundaries is at most r'.*

PROOF. Due to lemma 8 there is a mapping of the foreground components of A to the foreground components of $\hat{A}$. Lemma 7 states that this mapping is injective, and from lemmas 9 and 10 follows surjectivity. The same holds for the background components of A and $\hat{A}$. Due to lemma 10, both the foreground and background components of $\hat{A}$ are connected via direct pixel neighborhood. This implies a one-to-one mapping between the boundaries of A and $\hat{A}$, and thus identity of the homotopy trees.

It remains to be shown that for any dimension the Hausdorff distance between the boundaries of A and $\hat{A}$ is bounded by r'. Suppose to the contrary that $\partial\hat{A}$ contains a point $\vec{s}$ whose distance from ∂A exceeds r'. Due to the definition of an r'-grid, the sampling points of all pixels containing $\vec{s}$ are located within a hyper-sphere with center $\vec{s}$ and radius r'. Under the supposition, the inside ball of this hyper-sphere would either be completely inside or outside A, and the pixels were all either in $\hat{A}$ or $\hat{A}^c$. Thus, $\vec{s}$ could not be on $\partial\hat{A}'$ – contradiction. Therefore, the Hausdorff distance between ∂A and $\partial\hat{A}$ is at most r'. $\square$

This *geometric sampling theorem* does not only apply to regular, but also to irregular grids of any dimension, see fig. 5. Moreover, if the conclusion holds for a particular grid, it also holds for any translated and rotated copy of that grid.

However, the theorem does not say anything about topological equivalence between the original set and its reconstruction. The following theorem shows that the topology is preserved in 2D:

Theorem 12 (sampling theorem for 2-dimensional binary images) *A 2-dimensional r-regular shape $A \subset \mathbb{R}^2$ is weakly reconstructible with any r'-grid S in $\mathbb{R}^2$, $0 < r' < r$.*

PROOF. From theorem 11 we know that there is a one-to-one mapping between the boundaries of A and $\hat{A}$. Due to lemma 10, both the foreground and background components of $\hat{A}$ are connected via direct pixel neighborhood. In 2D this implies that the components of $\hat{A}$ are jordan curves. Since this is also the case for the boundary components of the original set A, an R^2-homeomorphism can be constructed between these boundaries, and A and $\hat{A}$ are R^2-topologically equivalent. Together with the statement of theorem 11 it follows that $\hat{A}$ is a weak reconstruction of A. $\square$

An even stronger result for the 2D case will be demonstrated in the next section. However, this theorem cannot be proved for higher dimensional spaces. To the contrary, it can be shown by constructing a counter example, that there exist shapes whose reconstruction is not topologically aquivalent to the original for some grids. We give the theorem for 3D, but its formulation for any dimension is straight-forward.

Theorem 13 *There exist 3D r-regular sets where one can find, for any $r' < r$, an r'-grid S such that the S-reconstruction is not topologically equivalent to the original.*

PROOF. Let $A \subset \mathbb{R}^3$ be two semi-spheres with radius r that are connected by a cylinder of the same radius. This shape is obviously r-regular. Consider a tangential plane P of the cylinder and a square P_S with radius r' lying in P and oriented so that one diagonal coincides with the line where P touches the cylinder, and the other is perpendicular to this direction, i.e. lies in the plane where the cylinder is curved. Translate P_S towards the center of the cylinder by an arbirary small ε. Then two opposite corners of the square are in A, while the other two are in A^c. Suppose the corners of the square are sampling points of a regular 3D square grid of suitable size and position (see fig. 6 left). Then the reconstruction $\hat{A}$ with respect to this grid contains pixels that are only indirectly connected to each other via a common edge, but not directly via a common face (6 right). Consequently, the surface of $\hat{A}$ is not a manifold, while the surface of A is. Therefore, the two surfaces cannot be topologically equivalent. $\square$

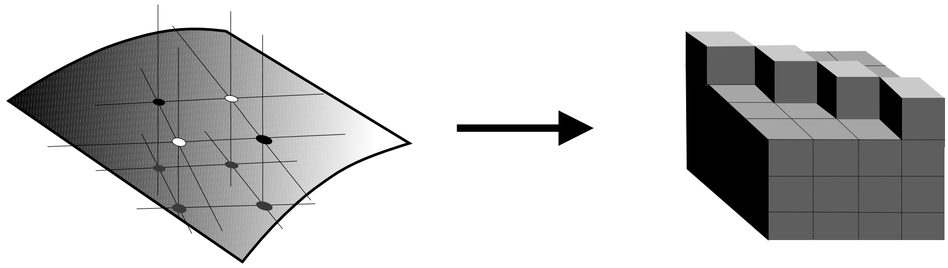

Figure 6. Construction of an r-regular counterexample: Given an r-regular cylinder, four sampling points on a square with radius $r' < r$ can be placed so that two diagonally opposite corners (black) are inside A and the two others (white) outside (left). The surface of the resulting reconstruction is not a manifold (right).

Obviously this example can be generalized to any dimension greater than three by adding more dimensions along which the curvature is zero. Thus r-regularity is not a sufficient criterion for shape preserving digitization in three or more dimensions in a strict topological sense when the reconstruction method according to definition 5 is used. Neither strong nor weak r-similarity can be guaranteed. It might thus seem that r-regularity is not a sufficient criterium for shape preservation in higher dimensions. However, we can hardly imagine a stronger criterion that still comprises a useful set of practically relevant shapes – r-regularity is already a very strong constraint. Therefore we believe that one should rather use alternative reconstruction methods (perhaps similar to the marching cubes algorithm) which guarantee that the reconstructed surfaces are manifolds. This problem must be further investigated.

7 Improved Sampling Theorems for 2D

The 2D sampling theorem we derived in the previous section is a generalization of the known theorems of Pavlidis [5] and Serra [7], because the 2D square and hexagonal grids on the one hand and topology preservation and identity of homotopy trees on the other hand are special cases of our definitions. So the results of Serra and Pavlidis are direct corollaries of theorem 12:

Corollary 14 *Let $S_1 := h_1 \cdot \mathbb{Z}^2$ be the square grid with grid size (minimal sampling point distance) h_1. Then every 2-dimensional r-regular set with $r > \frac{h_1}{\sqrt{2}}$ is weakly reconstructible with S_1. Let S_2 be the hexagonal grid with grid size h_2. Then every 2-dimensional r-regular set with $r > \frac{h_2}{\sqrt{3}}$ is weakly reconstructible with S_2.*

In case of Serra's theorem our proof even provides a better bound for the required sample distance (r vs. $r/\sqrt{3}$). But in 2D we can prove an even more powerful theorem by showing that (i) even strong r-similarity can be guaranteed, and (ii) r-regularity is not only a sufficient but also a necessary condition

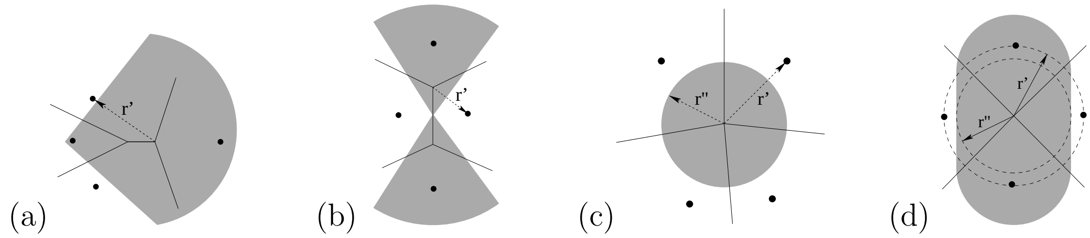

Figure 7. Examples where the topology of the reconstruction by an r-grid (balls and lines) differs from the topology of the original set (gray area) because it is not r-regular.

for a set to be reconstructible.

We first deal with the second property, which means that if A is not r-regular for some r, there exists an r'-grid S ($r' < r$) such that the S-reconstruction is not weakly r-similar to A. We conjecture, that this also holds for higher dimensions, but we could not yet prove the general case.

Theorem 15 *Let $A \in \mathbb{R}^2$ 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. 7 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 $\vec{p}_0$ that touches ∂A in at least two points. Draw a circle with radius r' around $\vec{p}_0$. *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. 7 c and d). $\square$

In the 2D case it is also possible to tighten sampling theorem 11 to strong r-similarity (recall that reconstructibility implies strong r-similarity, see definition 6).

Theorem 16 (strong sampling theorem for binary 2D images) *Let $r \in \mathbb{R}_+$ and $A \in \mathbb{R}^2$ an r-regular set. Then A is reconstructible with* any *r'-grid S, $0 < r' < r$.*

PROOF. From theorem 11 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 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 $\hat{\vec{c}}_i$ of $\partial \hat{A}'$ as follows (see fig. 8a):

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.

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. 8b). Let the special split point $\hat{\vec{c}}_i \in A$ (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 $\hat{\vec{c}}_i$. We must avoid that an extended part of $\partial A'$ gets mapped onto the single point $\hat{\vec{c}}_i$. Thus, we change the initial partitioning: Replace $\hat{\vec{c}}_i$ with two new split points $\hat{\vec{c}}_i{}'$ and $\hat{\vec{c}}_i{}''$, lying on $\partial \hat{A}'$ to either side of $\hat{\vec{c}}_i$ at a distance ε. Define as their corresponding split points the points $\vec{c}_i{}'$ and $\vec{c}_i{}''$ 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, $|\overline{\vec{c}_i{}'\hat{\vec{c}}_i}| < r'$ and $|\overline{\vec{c}_i{}''\hat{\vec{c}}_i}| < r'$, and the same is true for all points between $\vec{c}_i{}'$ and $\vec{c}_i{}''$. Therefore, ε can always be chosen so that every point between $\vec{c}_i{}'$ and $\vec{c}_i{}''$ can be mapped onto every point between $\hat{\vec{c}}_i{}'$ and $\hat{\vec{c}}_i{}''$ with a displacement of

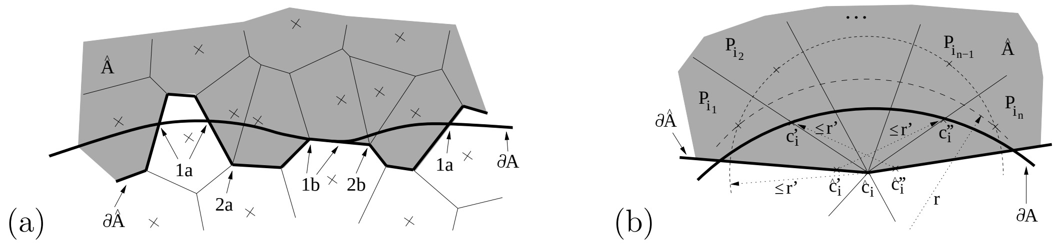

Figure 8. (a) Different cases for the definition of split points; (b) Partition refinement and mapping for case 2b.

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. 9a). 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 $\vec{s}_i$. If this point is not on $\partial A'$ (fig. 9b), C_i splits P_i into two parts, one containing $\hat{C}_i$ and the other containing $\vec{s}_i$. We now place r'-Reuleaux triangles as follows: a corner of every triangle is located at $\vec{s}_i$, 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 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 $\vec{s}_i$ of P_i is on $\partial A'$ (fig. 9c), this Reuleaux construction does not generally work. In this case, we first place two r'-Reuleaux triangles such that both have $\vec{s}_i$ as a corner point, one contains the start points $\vec{c}_s, \hat{\vec{c}}_s$ of C_i and $\hat{C}_i$ respectively, the other the end points $\vec{c}_e, \hat{\vec{c}}_e$, and they overlap $\hat{C}_i$ as much as possible. If they cover $\hat{C}_i$ completely, the Reuleaux construction still works with $\vec{s}_i$ 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 $\vec{s}_i'$ on the opposite normal with distance ε from $\vec{s}_i$ and project each point $\vec{c}$ of $\hat{C}_i$ not covered by either triangle onto the point where the line $\overline{\vec{c}\,\vec{s}_i'}$ crosses C_i. It can be seen that this mapping is an r'-homeomorphism: Draw circles with radius ε and $r' + \varepsilon$ around $\vec{s}_i'$. 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. □

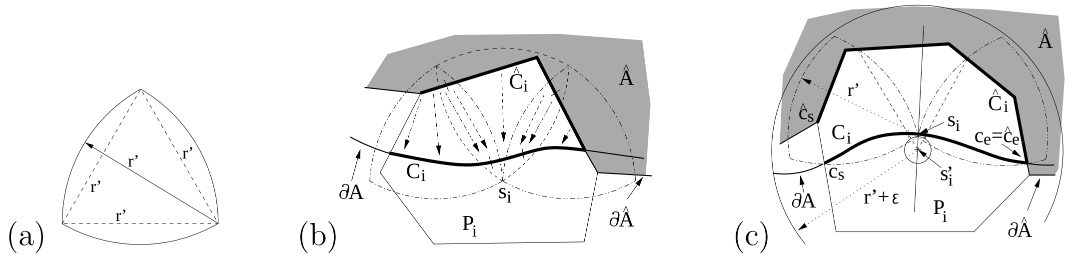

Figure 9. (a) Any two points in a Reuleaux triangle of size r' have a distance of at most r'; (b) Covering of corresponding segments with Reuleaux triangles; (c) Construction for sampling points lying on $\partial A'$.

In this section we proved two refinements of our sampling theorem, but the proofs rely on special properties of the 2-dimensional space. our results show the fundamentel role of r-regularity for shape preserving digitization:

Corollary 17 *Let $r \in \mathbb{R}_+$ and $A \subset \mathbb{R}^2$ Then A is r-regular if and only if A is reconstructible with* any *r'-grid S in $\mathbb{R}^2$ with $0 < r' < r$.*

From the last section we know that this cannot be directly generalized to higher dimensions. So the question remains open, how the class of sets, being reconstructible with any r'-grid with $r' < r$, looks like in higher dimensions.

8 Sampling of Blurred 2D Images

The subset digitization, which we used in the previous section, has the great disadvantage that it can not be realized in practice. In real imaging devices the analog image is always subjected to blurring before being digitized. The blurring has various causes. First, the light is diffracted in the lens. Second, there is defocus blur for objects that are not exactly in focus. Third, real sensor elements are not 0-dimensional but rather have finite area. The net effect of these factors can be modeled by a convolution of the analog image with a point spread function (PSF). Thus, instead of a binary image, a grayscale image is observed. In order to recover a binary reconstruction one can consider a particular level set $L_l = \{\vec{x} \in \mathbb{R}^2 | \hat{f}(\vec{x}) \geq l\}$ of the blurred image $\hat{f}$, i.e. apply a threshold. Since thresholding and digitization commute, we can perform the steps in a mathematically more convenient order by doing thresholding first and then digitizing the resulting level set by standard subset digitization, as can be seen in fig. 10. This order facilitates the following proofs, because we only have to analyse the properties of the level sets of an r-regular image and apply the already proved sampling theorem to them. So we begin with an analysis of the relationship between the original set and a level set of its blurred version, and then between the level set and its S-reconstruction. In order to get definitive results, we restrict ourselves to a particular type of PSF,

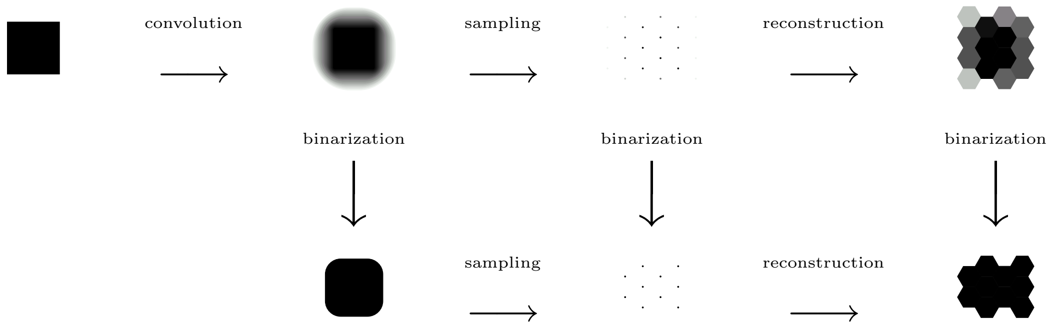

Figure 10. Model of digitization by a real camera (top row). Since binarization commutes with both sampling and reconstruction there are several mathematically equivalent ways from a blurred image to its discrete reconstruction.

namely flat disks of radius p. Although not very realistic, these PSFs have the advantage that the result of the convolution can be calculated by measuring the area of sets. In the sequel, A shall be a 2-dimensional r-regular set and k_p a disk-PSF with radius $p < r$. If $K_p(\vec{c})$ denotes the PSF's support region after translation to the point $\vec{c}$, the result of the convolution at $\vec{c}$ is given by:

$$\hat{f}(\vec{c}) = (k_p \star \chi_A)(\vec{c}) = \frac{\|K_p(\vec{c}) \cap A\|}{\|K_p(\vec{c})\|}$$

where $\star$ denotes convolution and $\|.\|$ is the area. Therefore, it is possible to derive properties of level sets by purely geometrical means. Obviously, all interesting effects occur in a $2p$-wide strip $A_p = \partial A \oplus K_p$ around the boundary ∂A, because out of this strip the kernel does not overlap ∂A, and the gray values are either 0 or 1 there ($\oplus$ denotes morphological dilation). Level sets have the following property:

Lemma 18 *Let $\vec{s}$ be a point on ∂A, and let $\vec{c}_1$ and $\vec{c}_2$ be the centers of the inside and outside osculating circles of radius r. Moreover, let $\vec{c}_3$ and $\vec{c}_4$ be the two points on the normal $\overline{\vec{c}_1\vec{c}_2}$ with distance p from $\vec{s}$. Then the boundary of every level set has exactly one point in common with $\overline{\vec{c}_3\vec{c}_4}$.*

PROOF. Consider a point $\vec{c}$ in $K_p(\vec{c}_3)$ and translate the line segment $\overline{\vec{c}_3\vec{c}_4}$ by $\vec{c} - \vec{c}_3$ (see fig. 11 left). Because of the restricted curvature of ∂A, the translated line segment intersects ∂A at exactly one point. Thus, as $t \in [0, 1]$ increases, the area of $\|K_p(\vec{c}_3 + t \cdot (\vec{c}_4 - \vec{c}_3)) \cap A\|$ is strictly decreasing. This area is proportional to the result of the convolution, so the same holds for the gray values. Since the p-ball centered in $\vec{c}_3$ is an inside osculating ball of A, the gray value at $\vec{c}_3$ is $f(0) = 1$. Likewise, $f(1) = 0$. This implies the lemma. $\square$

The curvature of the level set contours is bounded by the following lemma:

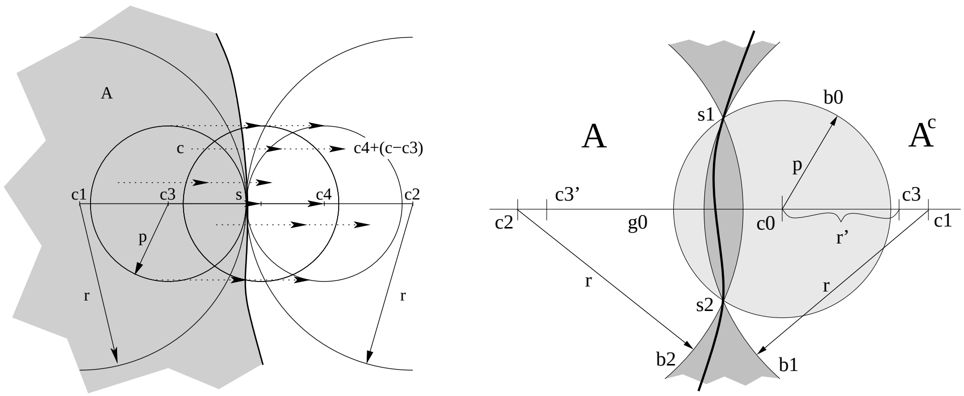

Figure 11. Left: If a p-ball is shifted orthogonally to the boundary ∂A from an inner osculating to an outer osculating position, its intersection area with A strictly decreases. Right: The boundary of the circle b_0 centered at point $\vec{c}_0$ (light gray) intersects the boundary of the set A (bold line) at the two points $\vec{s}_1$ and $\vec{s}_2$. Since A is r-regular, its boundary can only lie within the area marked with dark gray.

Lemma 19 *Let $\vec{c}_0 \in A_p$ be a point such that $(A \star K_p)(\vec{c}_0) = l$, $(0 < l < 1)$. Thus, $\vec{c}_0$ is part of level set L_l. Then there exists a circle b_{out} of radius $r_o \geq r' = r - p$ that touches $\vec{c}_0$ but is otherwise completely outside of L_l. Likewise, there is a circle b_{in} with radius $r_i \geq r'$ that is completely within L_l.*

PROOF. Consider the set $b_0 = K_p(\vec{c}_0)$ centered at $\vec{c}_0$. Let its boundary $\partial K_p(\vec{c}_0)$ intersect the boundary ∂A at the points $\vec{s}_1$ and $\vec{s}_2$ (see fig. 11 right). Let g_0 be the bisector of the line $\overline{\vec{s}_1\vec{s}_2}$. By construction, $\vec{c}_0$ is on g_0. Define $\vec{c}_1$ and $\vec{c}_2$ as the points on g_0 whose distance from $\vec{s}_1$ and $\vec{s}_2$ is r, and draw the circles b_1 and b_2 with radius r around them. Now, the boundary of A cannot lie inside either $b_1 \setminus b_2$ or $b_2 \setminus b_1$, because otherwise A could not be r-regular. The areas where ∂A may run are marked dark gray in fig. 11 right. Since $p < r$, there can be no further intersections between $\partial K_p(\vec{c}_0)$ and ∂A besides $\vec{s}_1$ and $\vec{s}_2$. On g_0, mark the points $\vec{c}_3$ between $\vec{c}_0$ and $\vec{c}_1$, and $\vec{c}_3'$ between $\vec{c}_0$ and $\vec{c}_2$, such that $|\overline{\vec{c}_1\vec{c}_3}| = |\overline{\vec{c}_2\vec{c}_3'}|$ and $\min(|\overline{\vec{c}_0\vec{c}_3}|, |\overline{\vec{c}_0\vec{c}_3'}|) = r' = r - p$. Due to the triangle inequality, and since $p < r$, such a configuration always exists. We prove the lemma for the circle b_{out} around $\vec{c}_3$, b_{in} around $\vec{c}_3'$ is treated analogously.

Let $b_3 = b_{\text{out}}$ be the circle around $\vec{c}_3$ with radius r', and b_3' the circle around $\vec{c}_3$ that touches $\vec{s}_1$ and $\vec{s}_2$ (fig. 12 left). Consider a point $\vec{c}_4$ on ∂b_3 and draw the circle b_4 with radius p around $\vec{c}_4$. This circle corresponds to the footprint of the PSF centered at $\vec{c}_4$. Now we would like to compare the result of the convolution $k_p \star \chi_A$ at $\vec{c}_0$ and $\vec{c}_4$. The convolution results are determined by the amount of overlap between A and $b_0 = K_p(\vec{c}_0)$ and $b_4 = K_p(\vec{c}_4)$ respectively. To compare $b_0 \cap A$ and $b_4 \cap A$, we split the two circles into subsets according to fig. 12 center (only b_0, b_4 and b_3' are shown in this figure). Circle b_0 consists of the subsets f_1, f_2, f_3, f_4, whereas b_4 consists of f_1, f_2, f_3', f_4'. The subsets f_1 and f_2 are shared by both circles, while due to symmetry f_3, f_3' and f_4, f_4'

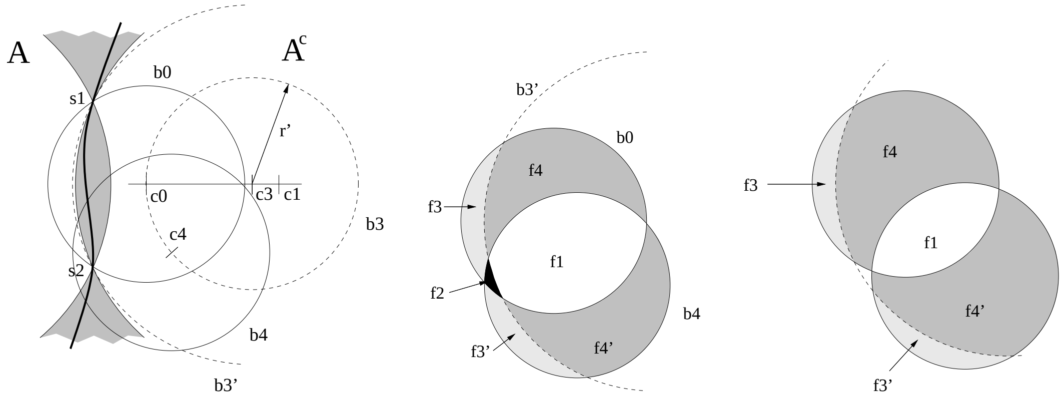

Figure 12. Left: The gray level at any point $\vec{c}_4 \neq \vec{c}_0$ on b_3 is smaller than the gray level at $\vec{c}_0$; center and right: decomposition of the circles b_0 and b_4 into subsets (see text).

are mirror images of each other. In terms of the subsets, we can express the convolution results as follows:

$$(k_p \star \chi_A)(\vec{c}_0) = \frac{\|f_1 \cap A\| + \|f_2 \cap A\| + \|f_3 \cap A\| + \|f_4 \cap A\|}{\|K_p\|}$$

$$(k_p \star \chi_A)(\vec{c}_4) = \frac{\|f_1 \cap A\| + \|f_2 \cap A\| + \|f_3' \cap A\| + \|f_4' \cap A\|}{\|K_p\|}$$

By straightforward algebraic manipulation we get:

$$\begin{aligned} \|K_p\| \left((k_p \star \chi_A)(\vec{c}_0) - (k_p \star \chi_A)(\vec{c}_4)\right) = \\ \|f_3 \cap A\| - \|f_3' \cap A\| + \|f_4 \cap A\| - \|f_4' \cap A\| \end{aligned} \tag{1}$$

Since the radius of b_3' is smaller than r, and its center $\vec{c}_3$ is between $\vec{c}_0$ and $\vec{c}_1$, the boundary $\partial b_3'$ intersects ∂A only at s_1 and s_2. It follows that subset f_3 is completely inside of A, whereas f_4' is completely outside of A. Hence, we have $\|f_3 \cap A\| = \|f_3\| = \|f_3'\|$ and $\|f_4' \cap A\| = 0$. Inserting this into (1), we get

$$\|K_p\| \left((k_p \star \chi_A)(\vec{c}_0) - (k_p \star \chi_A)(\vec{c}_4)\right) = \|f_3'\| - \|f_3' \cap A\| + \|f_4 \cap A\| > 0 \tag{2}$$

Thus, the gray level at $\vec{c}_4$ is smaller than l. When $\vec{c}_4$ is moved further away from $\vec{c}_0$, the subset f_2 will eventually disappear from the configuration (fig. 12 right). If $\vec{c}_3$ is outside of b_0, f_1 will finally disappear as well. It can easily be checked that (2) remains valid in either case. Due to the definition of $\vec{c}_3$, no other configurations are possible. Therefore, the gray values on the boundary ∂b_{out} are below l everywhere except at $\vec{c}_0$.

It remains to prove the same for the interior of b_{out}. Suppose the gray level at point $\vec{c} \in b^0_{\text{out}}$ were $l' \geq l$. By what we have already shown, the associated level line $\partial L_l'$ cannot cross the boundary ∂b_{out} (except at the single point c_0 if $l' = l$). So it must form a closed curve within b_{out}. However, this curve would cross some normal of ∂A twice, in contradiction to lemma 18. This implies the claim for outside circles. The proof for inside circles proceeds analogously. $\square$

We conclude that the shape of the level sets L_l is quite restricted:

Theorem 20 *Let A be an r-regular set, and L_l any level set of $k_p \star \chi_A$, where k_p is a flat disk-like point spread function with radius $p < r$. Then L_l is r'-regular (with $r' = r - p$) and strongly p-similar to A.*

PROOF. The proof of r'-regularity follows directly from the definition of r-regularity and lemma 19. The required p-homeomorphism can be constructed as follows: Because of the restricted curvature of ∂A, the normals of ∂A cannot intersect within the p-strip A_p around ∂A (cf. [3,4]). Therefore, due to 18, every point $\vec{s}$ on ∂A can be translated along its normal towards a unique point on the given level line ∂L_l and vice versa. The distance between s and its image is $\leq p$. This mapping can be extended to the entire $\mathbb{R}^2$-plane in the usual way, so that we get a p-homeomorphism with the desired properties. $\square$

This finally allows us to show what happens during the digitization of a set A that was subjected to blurring with a PSF:

Theorem 21 (sampling theorem for blurred binary 2D images)
Let A be an r-regular set, L_l any level set of $k_p \star \chi_A$, where k_p is a flat disk-like point spread function with radius $p < r$, and S a grid with maximum pixel radius $r'' < r - p$. The S-reconstruction $\hat{L}_l$ of L_l is strongly $(p + r'')$-similar to A.

PROOF. By theorem 20, L_l is r'-regular and there exists a p-homeomorphism between A and L_l. By theorem 16, the S-reconstruction of an r'-regular set with an r''-grid ($r'' < r'$) is strongly r''-similar to the original set. In other words L_l and $\hat{L}_l$ are r''-homeomorphic. Consequently, there is at least a $(p+r'')$-homeomorphism between A and $\hat{L}_l$. $\square$

Corollary 22 *Since $r'' + p < r$, any S-reconstruction of L_l is strongly r-similar to A, regardless of how the grid is rotated and translated relative to A and how the threshold is chosen.*

9 Conclusions

In this paper we proved a powerful geometric sampling theorem. Put simply, it means the following: When an r-regular set in an arbitrary dimensional space is digitized, the number and the inclusion properties of connected components of the set and its complement are preserved, and the digital components are directly connected. Parts that were originally connected do not get separated, and the Hausdorff distance between the original and reconstructed boundaries is at most half the pixel diameter. As these claims hold for any regular or irregular r'-grid in any dimension, they also hold for translated and rotated versions of some given grid. Thus, reconstruction is robust under Euclidean transformations of the grid or the shape. In 2D, an even stronger version of this theorem holds: Here, r-regularity is not only a sufficient, but also a necessary condition. Moreover, topology is preserved and boundary points are guaranteed to move at most half the pixel diameter, which ensures even better shape preservation.

The extension of our theorem towards blurred images also leeds to an intuitively very appealing result: When we digitize an ideal 2D binary image with *any* r''-grid, we can properly reconstruct a shape if it is r'-regular with $r' > r''$. But when the image is first subjected to blurring with a PSF of radius p, the set must be r-regular with $r > r'' + p$. In other words, the radius of the PSF must be added to the radius of the grid pixels to determine the regularity requirements for the original shape. It should also be noted that $r > r'' + p$ is a tight bound, which for instance would be reached if A consisted of a circle of radius r, and the threshold was 1 – in this case, any smaller circle could get lost in the reconstruction. However, for a single, pre-selected threshold a better bound can be derived.

The 2D sampling theorem 21 is closely related to the findings of Latecki et al. [3,4] about v-digitization (and thus also square subset and intersection digitization). In their approach, the grid must be square with sampling distance h, and the PSF is an axis aligned flat square with the same size as the pixels. Then, the pixel and PSF radius are both $r'' = p = h/\sqrt{2}$, and the original shape must be r-regular with $r > r'' + p = \sqrt{2}h$. This is exactly the same formula as in our case. We conjecture that our results can be generalized to a much wider class of radially symmetric PSFs, but we can't prove this yet.

Since strong r-similarity is also a local property, we can still apply our results if a 2D set is not r-regular in its entirety. We call a segment of a set's boundary *locally r-regular* if definition 3 holds at every point of the segment. Theorem 11 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 should be extended in several ways: First, the generalization to arbitrary dimensional spaces should be completed by proving or disproving the conjectures put forward in the paper. In particular, it would be interesting to know which sets can be reconstructed under preservation of strong r-similarity. Second, the results on blurred images should not only be extended towards higher dimensions, but also to more realistic point spread functions in order to be applicable to real cameras. It would also be useful to investigate what happens if the original shape was not uniformely colored but shaded, as would be the case for projections of real 3D scenes. Third, since r-regularity of shapes cannot always be guaranteed in practice (e.g. in the neighborhood of a junction) it is necessary to understand how the reconstruction may differ from the original and how this may be recognized and/or repaired.

References

[1] Giraldo, A., Gross, A., Latecki, L.J.: *Digitizations Preserving Shape.* Pattern Recognition **32**, pp. 365–376, 1999.

[2] Köthe U., Stelldinger, P.: *Shape Preserving Digitization of Ideal and Blurred Binary Shapes.* in: I. Nyström, G. Sanniti di Baja, S. Svensson (Eds.): Discrete Geometry for Computer Imagery, DGCI 2003, Lecture Notes in Computer Science 2886, pp. 82-91, Heidelberg: Springer, 2003.

[3] Latecki, L.J., Conrad, C., Gross, A.: *Preserving Topology by a Digitization Process.* Journal of Mathematical Imaging and Vision **8**, pp. 131–159, 1998.

[4] Latecki, L.J.: *Discrete Representation of Spatial Objects in Computer Vision.* Kluwer Academic Publishers, Dordrecht, 1998.

[5] Pavlidis, T.: *Algorithms for Graphics and Image Processing.* Computer Science Press, Rockville, Maryland. 1982.

[6] Ronse, C., Tajine, M.: *Discretization in Hausdorff Space.* Journal of Mathematical Imaging and Vision **12**, pp. 219–242, 2000.

[7] Serra, J.: *Image Analysis and Mathematical Morphology* Academic Press, New York, 1982.

[8] Stelldinger, P.: *Theoretische Grenzen der Bilddigitalisierung.* Diploma thesis, University of Hamburg, Department of Informatics, 2003.

[9] Stelldinger, P., Köthe U.: *Shape Preservation During Digitization: Tight Bounds Based on the Morphing Distance.* in: B. Michaelis, G. Krel (Eds.): Pattern Recognition, DAGM 2003, Lecture Notes in Computer Science 2781, pp. 108-115, Heidelberg: Springer, 2003.

[10] Tajine, M., Ronse, C.: *Topological properties of Hausdorff Discretizations.* ISMM'2000, Palo Alto (USA), pp. 41–50, Kluwer Academic Publishers, Dordrecht, 2000.