

Topologically Correct 3D Surface Reconstruction and Segmentation from Noisy Samples

Peer Stelldinger

University of Hamburg, 22527 Hamburg, Germany

Abstract. Existing theories on 3D surface reconstruction impose strong constraints on feasible object shapes and often require error-free measurements. Moreover these theories can often only be applied to binary segmentations, i.e. the separation of an object from its background. We use the Delaunay complex and α -shapes to prove that topologically correct segmentations can be obtained under much more realistic conditions. Our key assumption is that sampling points represent object boundaries with a certain maximum error. We use this in the context of digitization, i.e. for the reconstruction based on supercover and m -cell intersection samplings.

1 Introduction

A fundamental question of image analysis is how closely a computed image segmentation corresponds to the underlying real-world partition. Existing geometric sampling theorems are limited to binary partitions, where the space is split into (not necessarily connected) fore- and background components. In this case, the topology of the partition is preserved under various reconstruction schemes when the original regions are sufficiently smooth and the sampling is dense enough, e.g. see [1,2] for the case of 3D surface reconstruction.

However, these results have two important limitations: they do not make any predictions about the consequences of measurement errors, and they are not applicable when there are more regions than just fore- and background. While the second limitation is still valid today, there exist solutions for the first one: recently alternative surface reconstruction methods have been developed, which can deal with measurement errors [8,10].

Digital images consist of a finite number of sampling values in a regular grid. Segmentation means to group the sampling points (i.e. pixels) into meaningful regions. These regions can completely be described by their boundary. Thus segmentation can also be done by reconstructing the segment boundaries based on the subset of sampling points, which lie near the boundary. Our treatment of adaptively placed sampling points on the boundary is inspired by research on laser range scanning. Here, a number of isolated sampling points is scattered over the surface of the object of interest, and the task is to reconstruct the surface from the set of points. A successful solution of this problem is the concept of α -shapes [5,6]. The α -shape is essentially defined as the subset of the Delaunay

triangulation of the points where the Delaunay cells' radius is below $\alpha \in \mathbb{R}^+$. Under certain conditions, an α -shape is homotopy equivalent or even homeomorphic to the desired object surface. By applying this idea to the problem of image segmentation, a new condition on object shape could recently be derived that ensures homotopy equivalence of the digital segmentation with the original analog partitioning of the space [13]. In this work we prove such properties even when the segmentation is subject to measurement errors.

2 Preliminaries

To segment a geometric image means to partition the image space (i.e. the domain of the image function) into meaningful regions. The image space does not have to be the two-dimensional plane, e.g. for CT or MRT scans it is the three-dimensional space. Each region corresponds to a relevant (part of an) object in the real world and its reconstruction should preserve as much properties as possible. The partition of the image space to be recovered is defined as follows:

Definition 1. *Let the image space I be \mathbb{R}^n with $n \in \mathbb{N}$. A partition of the image space is defined by a finite set of pairwise disjoint regions $R = \{r_i \subset I\}$, such that each region $r_i \in R$ is a connected open set and the union of the closures of the regions covers the whole space, $\bigcup_i \bar{r}_i = I$. The boundary of the partition is $B := \bigcup_i \partial r_i$. Two regions r_i, r_j are called m -neighbors if the intersection $\bar{r}_i \cap \bar{r}_j$ contains an m -dimensional manifold with boundary, but no $(m+1)$ -dimensional manifold with boundary. Two $(n-1)$ -neighbors are also called direct neighbors.*

The simplest case of a partition is a *binary partition*, where the regions can be classified into foreground and background, such that every direct neighbor of a foreground region is a background region and vice versa. Then segmentation means separation of one (not necessarily connected) set from the background. Such a set is called a *shape*.

Most of the known results on topologically correct object or surface reconstruction are restricted to certain subclasses of shapes, having minimal bounds on the surface curvatures, like r -regular sets [9,4,12] or sets with certain local feature size [1,2]. This implies that regions cannot have corners, and junctions of three or more regions are impossible. These restrictions are somewhat relaxed by the notion of *r -halfregular partitions*, where an osculating r -ball must exist at least in the foreground or the background, and the topology must not change under either morphological opening or closing with a ball of radius $\leq r$ [11]. Corners are now possible, but the partition is still binary and has no junctions. In this paper, the class of feasible partitions of the space is extended as follows:

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

In other words, we can replace an infinitely thin boundary with a strip of width $2r$ such that the number and enclosure hierarchy of the resulting regions is

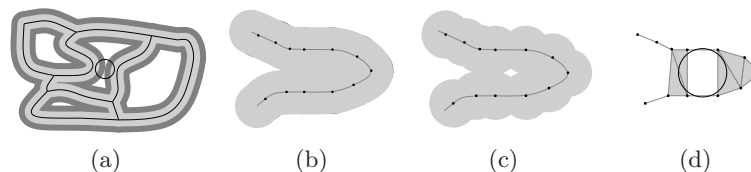


Fig. 1. (a) The homotopy type of an r -stable plane partition does not change when dilated with a disc of radius of at most r (light gray), while dilations with bigger radius (dark gray) may connect different parts as marked by the circle. The α -dilation (b) of the boundary of a two-dimensional α -stable partition may not be homotopy equivalent to the union (c) of the α -discs centered at the boundary sampling points. Thus the α -shape (d), which is always homotopy equivalent to the union of discs (c), may contain unwanted holes consisting of Delaunay triangles of radius greater than α . Thus there exists an α -disc centered in the hole which does not cover any boundary sampling point, as shown in (d).

preserved. In particular, “waists” are forbidden, whereas junctions are allowed, see Fig. 1(a). This includes r -regular and r -halfregular partitions, but also allows non-binary partitions and junctions and models real images much better. Since we want to deal with measurement errors (i.e. noise) when sampling the partition, we define a sampling of the surface as an approximation of the *boundary* of the partition with a finite set of *adaptively placed* sampling points. The sampling points are selected somehow “near” the boundary. We formalize this as follows:

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

Non-zero values of q can be caused by systematic or statistical measurement errors, but also by the sampling method used. Boundary sampling points may be determined in various ways (section 3), but this only matters in so far as it determines the accuracy of the sampling, i.e. the values of p and q . Once computed, we consider boundary sampling points as isolated points that somehow define the digital boundary and connect them by means of the Delaunay complex. Each element of the Delaunay complex is either the convex hull of a finite subset of the sampling points, such that all chosen sampling points lie on the boundary of a common hypersphere and no other sampling point is inside the hypersphere, or the intersection of two other elements of the complex. The hypersphere center of such a *Delaunay cell* is called the *center point* of the cell and the hypersphere radius is also called the *radius* of the cell. An m -dimensional cell of the complex is called *m-cell*. In order to approximate the boundary of the partition, we want to remove those edges and triangles from the Delaunay triangulation that are

not related to the boundary. A useful subset of the Delaunay complex is defined by the α -complex introduced in [5]:

Definition 4. *The α -complex $D_\alpha(S)$ of a set of points S is defined as the sub-complex of the Delaunay complex of S which contains all cells C such that*

- *the radius of the smallest sphere containing the sampling points of C is smaller than α , and it contains no other point of S , i.e. $C^0 \cap S = \emptyset$, or*
- *an incident cell C' with higher dimension is in $D_\alpha(S)$.*

The polytope $|D_\alpha(S)|$, i.e. the union of all elements of $D_\alpha(S)$, is called α -shape.

Since cells are removed from the Delaunay complex, the α -complex has holes which hopefully correspond to the regions we are trying to segment. In order to determine when this is the case, the following theorem is of fundamental importance (the proof can be found in [6]):

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

Consequently, the α -shape $|D_\alpha|$ is homotopy equivalent to the original partition of the space if and only if the dilation of the boundary sampling points with α -balls is homotopy equivalent to the boundary of the partition. This requirement is indeed fulfilled in certain situations: In [4] it is proved that $|D_\alpha|$ is even homeomorphic to B if B is the boundary of a two-dimensional r -regular set with $p < \alpha < r$ and $q = 0$. In three dimensions the authors recently derived an analog result [12]. There the α -shape itself cannot be guaranteed to be homeomorphic, but it can be used to derive a homeomorphic surface approximation in a very simple way: with defining the outer boundary of the α -shape as the union of all triangles of the corresponding α -complex, which can be seen from the outside (i.e. from a point being outside the original object), one gets the following theorem (the proof can be found in [12]):

Theorem 2 (Stelldinger). *Let A be a three-dimensional r -regular set and S be an α -sampling of its boundary ∂A such that $2\alpha < r$. Then the polytope $|D_\alpha|$ is of the same homotopy type as ∂A , and the outer boundary of $|D_\alpha|$ is homeomorphic to ∂A .*

Unfortunately, these approaches no longer apply when the original partition is not r -regular and/or the boundary sampling points are not exactly on the original boundary, i.e. they are noisy. Fig. 1(b)-(d) shows a two-dimensional example where the r -dilation of the boundary is homotopy equivalent to the boundary (i.e. the partition is r -stable), but the dilation of the boundary sampling points is not. This problem has already been solved for the two-dimensional case in [13]. There it is shown that by filling small regions one gets a boundary representation with correct homotopy type such that the separated regions are homeomorphic to the original ones. These small regions can uniquely be identified if the sampling is dense enough. In three or higher dimensional spaces filling such small regions is not enough, since reconstruction artifacts can have a more complicated topological structure, like e.g. tunnels or bridges. The rest of the paper is

devoted to the question what can be said under these more general conditions in case of higher dimensional spaces.

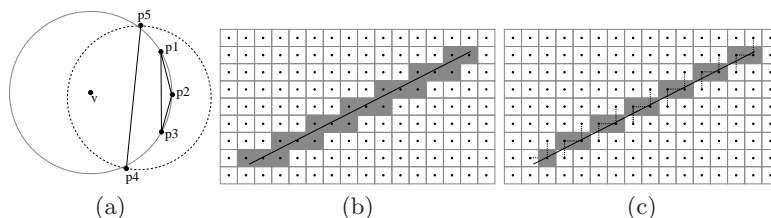


Fig. 2. (a) Any circumcircle around p_4 and p_5 contains p_1 , p_2 , and p_3 (see text). (b) The *supercover digitization* contains all sampling points whose pixel facets intersect the arc. (c) Where the boundary intersects the dual grid, the nearest sampling points form the *grid intersection digitization*.

As can be seen in Figure 1(b)-(d), if we have an object which is not r -regular, but also if the sampling is not strict, i.e. a (p, q) -sampling with $q > 0$, the complement of the α -shape reconstruction may have new small regions, which lie inside the α -dilation of the original boundary. In order to get a topologically correct boundary reconstruction we must at least fill these small regions. In the following, the components of $|D_\alpha(S)|^c$ will be called α -holes. As we will see, the spurious holes are restricted in their size. Thus we define (α, β) -holes in order to distinguish between spurious and wanted α -holes:

Definition 5. Let $D_\alpha(S)$ be the α -complex of a sampling S and $|D_\alpha(S)|$ be its α -shape. Then the α -holes of $|D_\alpha(S)|$ are the components of $|D_\alpha(S)|^c$. The (α, β) -holes of $|D_\alpha(S)|$ are the α -holes H , where the largest radius of some n -cell in H is at least $\beta \geq \alpha$. The union of the α -shape $|D_\alpha|$ with all α -holes of D_α that are not (α, β) -holes is called the (α, β) -shape reconstruction.

For simplicity, we also use the term “hole” for the component which contains the infinite region. It is an (α, β) -hole for arbitrary large β . It follows from Theorem 1 that there is a 1-to-1 relation between α -holes and the holes in the union of α -discs around the sampling points. The following lemma establishes that a similar relationship exists for (α, β) -holes (we prove the lemma for the n -dimensional case):

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

Proof. **I** ($d_H(v \in h, S) \geq \beta \Rightarrow h$ is an (α, β) -hole): when v is in the infinite region, the claim follows immediately. Otherwise, v is contained in some Delaunay n -cell. By assumption, the corners of this triangle must have distance $\geq \beta$ from v . Therefore, the radius of the n -cell must be at least β , and the claim follows. **II** (h is an (α, β) -hole $\Rightarrow \exists v \in h$ with $d_H(v, S) \geq \beta$): by assumption, the closure

of h contains a Delaunay n -cell t with radius of at least β . Consider its center point v (i.e. the center of its Delaunay sphere). If it is within the n -cell t , it is also in h and the claim follows. Otherwise, it is at least in some (α, β) -hole, and we must prove that t is in the same hole. Suppose to the contrary that v and t are in different α -holes. Then there exists a Delaunay cell t' (this does not have to be an n -cell) between t and v whose covering radius (i.e. the radius of the smallest covering ball) is smaller than α . The corners of t' cannot be inside the Delaunay sphere of t because otherwise t would not be a Delaunay n -cell. t' cannot contain v because its covering radius would then be at least β . Now consider Figure 2(a). It shows for the two-dimensional case the 2-cell t with corners p_1, p_2, p_3 and its Delaunay circle (gray) with center point v . The points p_4 and p_5 are the end points of one side of t' . Their distance $|p_4p_5|$ must be greater than $|p_1p_3|$. Consequently, any covering circle with radius $\leq \alpha$ (dashed) around p_4 and p_5 contains t , contrary to the condition that it must not contain any other sampling point. This obviously also holds in higher dimensions. The claim follows from the contradiction. \square

Now we can use the notion of (α, β) -holes to “repair” α -complexes that contain too many holes. After filling all α -holes which are not (α, β) -holes we get a one-to-one-mapping of the components of ∂A^c to the components of the complement of the (α, β) -shape reconstruction:

Theorem 3. *Let \mathcal{P} be an r -stable partition of the space \mathbb{R}^n , and S be a (p, q) -sampling of \mathcal{P} 's boundary B . Then the (α, β) -shape reconstruction \mathcal{R} preserves connectivity and neighborhood relations and defines a one-to-one-mapping of the (α, β) -holes of \mathcal{R} to the regions r_i of \mathcal{P} , if (1) $p < \alpha \leq r - q$, (2) $\beta = \alpha + p + q$ and (3) every region r_i contains an open γ -disc with $\gamma \geq \beta + q > 2(p + q)$.*

Proof. Let U be the union of open α -balls centered at the points of S . Furthermore, let $B^\oplus = B \oplus \mathcal{B}_{\alpha+q}^0$ be the dilation of B with an open $(\alpha + q)$ -ball, and $r_i^\ominus = r_i \ominus \mathcal{B}_{\alpha+q}$ the erosion of region $r_i \in \mathcal{P}$ with a closed $(\alpha + q)$ -ball.

- According to the definition of a (p, q) -sampling, the dilation of B with a closed q -ball covers S . Consequently, B^\oplus covers U . Therefore, U cannot have fewer connected components than B^\oplus . B^\oplus has as many components as B due to the r -stability of the partition \mathcal{P} . Conversely, since $\alpha > p$, every open α -ball around a point of S intersects B , and the union U of these balls covers the entire boundary B . It follows that U cannot have more components than B . The number of components of B and U is thus equal. Due to the same homotopy types of U and $|D_\alpha|$ (according to Lemma 1), this also holds for the components of $|D_\alpha|$.
- Since \mathcal{P} is r -stable with $r \geq \alpha + q$, each r_i^\ominus is a connected set with the same topology as r_i . The intersection $r_i^\ominus \cap B^\oplus$ is empty, and r_i^\ominus cannot intersect $U \subset B^\oplus$ and $|D_\alpha| \subset U$. Hence, r_i^\ominus is completely contained in a single α -hole of $|D_\alpha|$.
- Due to condition 3, r_i contains a point whose distance from B is at least $\gamma = \beta + q$. Its distance from S is therefore at least $\gamma - q = \beta$. Due to Lemma 1, the α -hole which contains r_i^\ominus is therefore also an (α, β) -hole.

- Since B^\oplus covers U and U covers B , no (α, β) -hole can intersect both r_i^\ominus and r_j^\ominus ($i \neq j$). It follows from this and the previous observation, that every region r_i can be mapped to exactly one (α, β) -hole which will be denoted h_i .
- An α -hole that does not intersect any region r_i^\ominus must be completely contained within B^\oplus . Every point $v \in B^\oplus$ has a distance $d < \alpha + q$ to the nearest point of B . In turn, every point in B has a distance of at most p to the nearest point in S . Hence, the distance from v to the nearest point of S is $d' < \alpha + p + q = \beta$. According to Lemma 1, this means that an α -hole contained in B^\oplus cannot contain an n -cell with radius β and cannot be an (α, β) -hole.
- The previous observation has two consequences: (i) All holes remaining in \mathcal{R} intersect a region r_i^\ominus . Therefore, the correspondence between r_i and h_i is 1-to-1, and B and $|\mathcal{R}|$ enclose the same number of regions. (ii) All differences between \mathcal{R} and D_α (i.e. all Delaunay cells re-inserted into \mathcal{R}) are confined within B^\oplus . This implies that $|\mathcal{R}|$ cannot have fewer components than B^\oplus and B . Since all re-inserted cells are incident to D_α , $|\mathcal{R}|$ cannot have more components than $|D_\alpha|$, which has as many components as B (see first observation). Hence, B and $|\mathcal{R}|$ have the same number of components.
- Consider the components of the complement $(r_i^\ominus)^C$ and recall that r_i^\ominus is a subset of both r_i and h_i for any i . Since B and $|\mathcal{R}|$ have the same number of components, it is impossible for h_i^C to contain a cell that connects two components of $(r_i^\ominus)^C$. This means that the sets r_i^C and h_i^C have the same number of components. This finally proves that the constructed one-to-one-mapping preserves the neighborhood relations. \square

Filling spurious holes in the α -shape reconstruction is a necessary step for getting a topologically correct boundary reconstruction. But for $n \geq 3$ there are also other problems regarding topology: although the (α, β) -shape reconstruction separates the different regions from each other, these regions may have small tunnels and/or other topological changes inside B^\oplus . In order to identify and remove these cases, we will at first apply a homotopy type preserving thinning:

We will denote an m -cell C in a cell complex D as *simple* if the number of cells of D which contain C is equal to one. Now the containing cell must be an $(m + 1)$ -cell and the removal of the two cells does neither change the homotopy type of the complex nor the topology of the background regions. Now the thinning algorithm for the (α, β) -shape reconstruction is as follows:

1. Find all simple m -cells ($n > m \geq 0$) of the given (α, β) -shape reconstruction and put them in a priority queue (the sorting will be discussed below).
2. As long as the queue is not empty:
 - (a) Get the m -cell e with the highest priority from the queue.
 - (b) If e is not simple anymore, it has lost this property during the removal of other cells. Skip the following and recommence with step 2.
 - (c) Otherwise, remove e and the adjacent $(m + 1)$ -cell $t \in \mathcal{R}$ from the boundary reconstruction.
 - (d) Check whether the other cells adjacent to t have now become simple and put them in the queue if this is the case.

Obviously the algorithm terminates for any finite cell complex, and the resulting boundary reconstruction contains no n -cells, i.e. it is thin. Since we want a boundary reconstruction which is as simple as possible in a topological sense (i.e. as few as possible tunnels, etc.), but which still separates the different regions from each other, we want to have a cell complex, where every cell is adjacent to at least two different background components (i.e. regions of the (α, β) -shape reconstruction). Thus we remove every cell, which is not adjacent to two different background regions. Since an $(n - 1)$ -cell, which is not adjacent to two different background regions, will already be removed by the above thinning algorithm, we only have to check m -cells with $m < n - 1$. This can be done locally, since these cells are characterized by having no adjacent $(m + 1)$ -cell in the complex. Thus the whole algorithm is as follows:

Thinned (α, β) -shape reconstruction algorithm:

1. Given a (p, q) -sampling S of the boundary of some partition of the space, compute the α -complex of S with some $\alpha > p$.
2. Add all cells to the complex, which belong to an α -hole which is no (α, β) -hole for $\beta = \alpha + p + q$.
3. Find all simple m -cells (for any m with $n > m \geq 0$) of the given (α, β) -shape reconstruction and put them in a priority queue (the sorting will be discussed below).
4. As long as the queue is not empty:
 - (a) Get the m -cell e with the highest priority from the queue.
 - (b) If e is not simple anymore, it has lost this property during the removal of other cells. Skip the following and recommence with step 4.
 - (c) Otherwise, remove e and the adjacent $(m + 1)$ -cell $t \in \mathcal{R}$ from the boundary reconstruction.
 - (d) Check whether the other cells adjacent to t have now become simple and put them in the queue if this is the case.
5. For m going from $n - 2$ to 0 do:
 - (a) Remove all m -cells of the complex, which do not have an adjacent $(m + 1)$ -cell in the complex.

Theorem 4. *Let \mathcal{P} be an r -stable partition of the space \mathbb{R}^n , and S be a (p, q) -sampling of \mathcal{P} 's boundary B . Then the thinned (α, β) -shape reconstruction algorithm results in a cell complex D with $|D|$ having the same homotopy type as B , and the components of B^c are topologically equivalent to the components of $|D|^c$, if (1) $p < \alpha \leq r - q$, (2) $\beta = \alpha + p + q$ and (3) every region r_i contains an open γ -disc with $\gamma \geq \beta + q > 2(p + q)$.*

Proof. The resulting reconstruction obviously separates the components of $|D|^c$ from each other, which can be mapped one-to one onto the components of B^c . Since the α -ball reconstruction covers B , the α -shape reconstruction contains a polygonal surface which is of the same homotopy type as B . Thus also the (α, β) -shape reconstruction contains such a polygonal surface. Applying only the thinning algorithm results in a an object B' being a deformation retract of

the (α, β) -shape reconstruction. Thus it also contains a polygonal surface B'' , which is of the same homotopy type as B . This surface B'' is everywhere thin, i.e. any of its m -cells with $m \leq n - 2$ has at least two neighboring $m + 1$ -cells in B'' . Thus B'' remains unchanged during the cell complex simplification. Any other cell of B' will be removed. Thus $B'' = |D|$. Since the boundaries of any component of B^c and the corresponding component of $|D|^c$ are not only of the same homotopy type but also homeomorphic (because both are composed of components of the same homotopy type, which are all $(n - 1)$ -dimensional manifolds without boundary), the components are also homeomorphic to each other. \square

The complexity of the algorithm is dominated by the Delaunay tetrahedrization which is known to be $O(n^2)$ in the number of sampling points. As far as region topology is concerned, the ordering of the m -cells in the priority queue is arbitrary. But we think, that orderings should be favored, which lead to visually appealing results, e.g. by emphasizing flat surfaces. This can be done in the following way:

Definition 6. *The minimal (α, β) -shape reconstruction is the result of the thinned (α, β) -shape reconstruction algorithm, when using radii of the simple m -cells as priority, i.e. m -cells with big radius are the first to be removed.*

When only using the thinning algorithm after (α, β) -shape reconstruction, the resulting regions are correctly separated from each other. Moreover by using the cell radii as ordering criterion, the resulting hypersurface is as smooth as possible, since the size of the $(n - 1)$ -cells is minimized. Thus, since a minimal boundary reconstruction is a shortest possible one with correct topology, the surviving edges connect sampling points closest to each other. Neighboring sampling points therefore align in an optimal way on the thinned boundary.

3 Application to Sampling Schemes

In Theorem 3, p and q are assumed given. We now make their meaning and consequences more intuitive, by computing them for two of the most common sampling schemes. Given a sampling grid S (e.g. a cubic grid), we want to define a subset of sampling points which approximates the boundaries of a partition of the space. Obviously simple subset sampling of the boundary is not a good choice, since in general hardly no sampling point will lie exactly on this surface. Alternatively one can choose the set of sampling points whose voxels (Voronoi regions) intersect the boundary of the partition. Such a method is called *supercover sampling*, since it is related to supercover digitization [7].

Definition 7. *Let $S \subset \mathbb{R}^n$ be an r' -grid (i.e. the maximal distance from any point to the nearest sampling point is at most r'). The supercover sampling of a set $A \subset \mathbb{R}^n$ based on S is the set S' of all sampling points $s \in S$ whose hypervoxel intersects A , see Figure 2(b).*

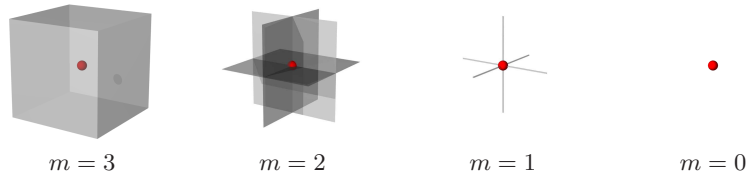


Fig. 3. Given a sampling point (s) in a cubic grid, the shaded region shows the intersection of $\mathcal{V}_S(s)$ with the union of all adjacent m' -cells of the Delaunay complex for $m' \leq m$.

Lemma 2. *The supercover sampling S' of a set A based on an r' -grid, is an (r', r') -sampling of A .*

Proof. The distance of any sampling point in S' to the nearest point in A can be at most r' , since S' is based on an r' -grid. Since the hypervoxels of the supercover sampling cover A , the distance of any point of A to the nearest sampling point in S is also at most r' . \square

If the object to be digitized is a curve in \mathbb{R}^n , the supercover sampling is $(n-1)$ -connected (i.e. connected via cells of dimension of at least $n-1$), since for any point on the curve there exists a neighborhood, which is covered by the hypervoxels. However, one often wants a curve to be represented by a sampling, which is “as thin as possible”, i.e. only 0-connected. This is fulfilled for the square grid in two dimensions by the grid intersection sampling [7], see Figure 2(c). It is well-known, that the grid intersection digitization is a subset of the supercover digitization on a square grid. The grid intersection sampling can easily be generalized to arbitrary grids in any dimension:

Definition 8. *Let $S \subset \mathbb{R}^n$ be a sampling grid. Further, for any sampling point $s \in S$ let $G_m(s)$ be the intersection of the hypervoxel $\mathcal{V}_S(s)$ with the union of all m' -cells, $m' \leq m$, of the Delaunay complex of S being adjacent to s . Then the m -cell intersection sampling of a set $A \subset \mathbb{R}^n$ is defined as the union S' of all sampling points $s \in S$, where $G_m(s) \cap A$ is not empty, see Figure 2(c) for an example of the 1-cell intersection sampling in 2D.*

Thus the 1-cell intersection sampling based on a square grid is equal to the grid intersection sampling. Moreover the 0-cell intersection sampling is the same as the subset sampling and the n -cell intersection sampling equals the supercover sampling in \mathbb{R}^n . This directly implies that an m_1 -cell intersection sampling is always subset of an m_2 -cell intersection sampling of a given object, if they are based on the same sampling grid and if $m_1 < m_2$. Figure 3 shows $G_m(s)$ for different m in case of a cubic grid. While the m -cell intersection sampling of a connected set does not need to be connected in case of $m < n-1$, it is $(n-1)$ -connected if $m = n$, since then it is equal to the supercover sampling. Moreover for $m = n-1$ it is 0-connected, if the grid is not degenerated:

Lemma 3. *Let S be the $(n - 1)$ -cell intersection sampling of a connected set based on a not degenerated grid. Then S is 0-connected.*

Proof. If A is empty or if A consists of only one point, the proof is obvious. Otherwise let x, y be two arbitrary points in A and let P be a path in A from x to y . If x, y and p lie in the same n -cell of the Delaunay complex, there is nothing to show. Thus let p go through at least two n -cells of the Delaunay complex. Since the union of the $(n - 1)$ -cells of the complex is equal to the union of all $G_{n-1}(s)$, and since any two sampling points with a common n -cell in the Delaunay complex are always at least 0-connected, the $(n - 1)$ -cell intersection sampling of p is 0-connected. Thus S must also be 0-connected. \square

Now we will show that the $(n - 1)$ -cell intersection sampling is of higher sampling accuracy (i.e. lower q) than the supercover sampling, while the sampling density (i.e. the smallest possible value of p) is not as high.

Lemma 4. *Let S' be a not degenerated r' -grid. When each component of a set A is intersected by at least one $(n-1)$ -cell of the Delaunay complex of S' , the $(n-1)$ -cell intersection digitization S of A based on the grid S' , is a $(2r', q)$ -sampling of A' with $q < r'$. If S' is a Cartesian r' -grid, S is even a $(2r', \frac{\sqrt{n-1}}{\sqrt{n}}r')$ -sampling.*

Proof. The set of $(n - 1)$ -cells the Delaunay complex of a not degenerated grid partitions the space, such that any n -cell has a diameter of at most $2r'$. Thus, since each component of A intersects at least one $(n - 1)$ -cell, the distance of any point of A to the nearest sampling point of S is at most $2r'$.

For any sampling point $s \in S$ there exists an adjacent $(n - 1)$ -cell which intersects A inside the hypervoxel of S . Any such intersection point has distance of smaller than r' to s . In case of a Cartesian grid, the sidelength of the $(n - 1)$ -cells is $\frac{2}{\sqrt{n}}r'$ and the largest distance of a point in the intersection of the hypervoxel and an $(n - 1)$ -cell of s is $\frac{\sqrt{n-1}}{\sqrt{n}}r'$. \square

We can now simply deal with sampling errors due to noise or blurring. We just have to add the expected positional error caused by these influences to the above computed positional error caused by the sampling.

4 Conclusions

This paper describes how to reconstruct a surface topologically correct from a sufficiently dense set of surface sampling points in the presence of measurement errors due to sampling but also due to noise and other influences. The theorem applies to a much wider class of shapes (r -stable partitions) than previous approaches. The situation in real images is thus modeled much more faithfully because shapes may now have corners and junctions, and standard segmentation algorithms can be used. Moreover an analysis on the amount of the measurement errors is given for the case of some of the mostly used digitization methods.




Fig. 4. Artificial boundary samples derived from a CT scan of the Stanford bunny (on courtesy of the Stanford Volume Data Archive graphics.stanford.edu/data/voldata/). Left: Sparse subset of boundary voxels due to 2-cell intersection sampling. Right: Result of thinned (α, β) -shape reconstruction.

References

1. Amenta, N., Bern, M. and Kamvysselis, M.: *A new Voronoi-based Surface Reconstruction Algorithm*, Proceedings of the 25th annual Conference on Computer Graphics and Interactive Techniques, 415–421 (1998).
2. Amenta, N., Choi, S., Dey, T.K., and Leekha, N.: *A Simple Algorithm for Homeomorphic Surface Reconstruction*, Proceedings of the 16th annual Symposium on Computational Geometry, 213–222 (2000).
3. Attali, D.: *r-Regular Shape Reconstruction from Unorganized Points*, Computational Geometry: Theory and Applications, Vol. 10, No. 4, 239–247 (1998).
4. Bernardini, F. and Bajaj, C.L.: *Sampling and Reconstructing Manifolds Using Alpha-Shapes*, Proc. 9th Canadian Conf. Computational Geometry (1997).
5. Edelsbrunner, H. and Mücke, E.P.: *Three-dimensional alpha shapes*, ACM Trans. Graphics, 13:43-72 (1994).
6. Edelsbrunner, H.: *The union of balls and its dual shape*, Discrete Comput. Geom., 13:415-440 (1995).
7. Klette, R. and Rosenfeld, A.: *Digital Geometry*, Morgan Kaufman (2004).
8. Kolluri, R., Shewchuk, J.R., and O’Brien, J.F.: *Spectral surface reconstruction from noisy point clouds*, Eurographics Symp. on Geom. Proc. (2004).
9. Latecki, L.J., Conrad, C., and Gross, A.: *Preserving Topology by a Digitization Process*, J. Mathematical Imaging and Vision, 8:131-159 (1998).
10. Mederos, B., Amenta, N., Velho, L., and de Figueiredo, L.H.: *Surface reconstruction from noisy point clouds*, Eurographics Symp. on Geom. Proc. (2005).
11. Stelldinger, P.: *Digitization of Non-regular Shapes*, Mathematical Morphology, Proc. of ISMM (2005).
12. Stelldinger, P.: *Topologically Correct Surface Reconstruction Using Alpha Shapes and Relations to Ball-Pivoting*, Submitted 2007.
13. Stelldinger, P., Koethe, U., and Meine, H.: *Topologically Correct Image Segmentation Using Alpha Shapes*, Proc. of DGCI (2006).