Voxel Cores: Efficient, robust, and provably good approximation of 3D medial axes
(Supplementary Proofs)

1 Nearest voxel vertex

Several proofs of the paper (Theorem 3.1, Lemma 3.6) rely on the assumption that the closest voxel vertex to a point on a voxel element (e.g., edge or face) must be a vertex of that element. Here we prove this assumption for voxel shapes in arbitrary dimensions.

Lemma 1.1. *Consider the n-dimensional integer grid $\mathbb{Z}^n$ and a point $x \in \mathbb{R}^n$. Let e be an element of the grid (e.g., vertex, edge, face, etc.) that contains x. Then the nearest grid vertices to x are vertices of e.*

Proof. It suffices to show that, for any grid vertex v, there exists some vertex v' of e such that $d(x, v') \leq d(x, v)$, where the equality holds only if $v = v'$ (see illustration below).

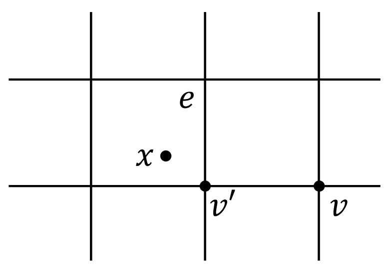

Figure 1: Notations used in the proof.

Let the dimensionality of e be $d \leq n$, and let D be the set of d coordinate axes that span the d-dimensional subspace containing e. Note that all points of e share a common (integer) coordinate along any of the remaining $n - d$ coordinate axes. Let the coordinates of x and v be $\{x_1, \ldots, x_n\}$ and $\{v_1, \ldots, v_n\}$, we define $v' = \{v'_1, \ldots, v'_n\}$ as follows:

$$v'_i = \begin{cases} x_i, & \text{if } i \notin D \\ \lceil x_i \rceil, & \text{if } i \in D \text{ and } v_i \geq x_i \\ \lfloor x_i \rfloor, & \text{if } i \in D \text{ and } v_i \leq x_i \end{cases} \quad (\forall i = 1, \ldots, n) \tag{1}$$

It is easy to see that v' is a vertex of e. Furthermore,

$$\begin{aligned} d(x,v)^2 &= \sum_{i \notin D} (x_i - v_i)^2 + \sum_{i \in D} (x_i - v_i)^2 && (2) \\ &\geq \sum_{i \in D} (x_i - v_i)^2 && (3) \\ &\geq \sum_{i \in D} (x_i - v'_i)^2 && (4) \\ &= \sum_{i \notin D} (x_i - v'_i)^2 + \sum_{i \in D} (x_i - v'_i)^2 && (5) \\ &= d(x,v')^2 && (6) \end{aligned}$$

where the equalities hold only if $v_i = x_i$ for $i \notin D$ (first equality) and $v_i = v'_i$ for $i \in D$ (second equality). In other words, $v = v'$. □

2 Properties of the voxel core

We will prove our main results on the voxel core, its homotopy (Theorem 3.4) and proximity (Theorem 3.5).

2.1 Homotopy

Recall that a voxel shape $\mathcal{O}$ is the interior of a union of voxels, and the voxel core $\mathcal{C}$ is the subset of Voronoi elements whose dual Delaunay elements in the Delaunay triangulation of the boundary vertices intersect $\mathcal{O}$.

Theorem 2.1 (Theorem 3.4 in the paper)**.** *$\mathcal{C}$ is homotopy equivalent to $\mathcal{O}$.*

Proof. The homotopy equivalence will be shown by examining a sequence of spaces. Consider the abstract cell complex, $\widehat{\mathcal{O}}$, built from the 3-dimensional Delaunay cells intersecting $\mathcal{O}$ (which are dual to vertices of $\mathcal{C}$) that are glued along their common two dimension faces (which are dual to edges of $\mathcal{C}$). Note that $\widehat{\mathcal{O}}$ is not the same as the closure of $\mathcal{O}$, which has additional identifications along boundary edges and vertices. From $\widehat{\mathcal{O}}$, we will build another cell decomposition, X, by removing a regular neighborhood of $\partial\widehat{\mathcal{O}}$ from $\widehat{\mathcal{O}}$. Note that since $\mathcal{O}$ is homoeomorphic to an open subset of Euclidean space, it is a manifold. The collar theorem then implies that $\mathcal{O}$ and X are homeomorphic [4].

As X is a cell complex, Björner's nerve theorem implies that X is homotopy equivalent to the nerve, N, of the 3-dimensional cells of X [1]. This nerve N is an abstract simplicial complex, with a vertex for each 3-dimensional cell of X and a k-simplex for any set of k 3-dimensional cells of X with common intersection. By our construction of X, N has one edge for each Delaunay face that intersects $\mathcal{O}$, and a k-simplex for each Delaunay edge e that intersects $\mathcal{O}$ and bounds k 3-dimensional Delaunay cells.

We claim that N is homotopy equivalent to $\mathcal{C}$. Note that $\mathcal{C}$ is *almost* a geometric realization of N, with the exception that for each 2-dimensional face in $\mathcal{C}$ with k sides, there is a k-simplex in N. Consider the map $f : N \to \mathcal{C}$ that sends each vertex of N to the corresponding vertex in C and extends linearly to the simplices of N. Notice that f projects each of the k-simplices of N to a k-gon in $\mathcal{C}$, and the preimage of any point on a k-gon in $\mathcal{C}$ is homeomorphic to the intersection of a hyperplane and a simplex of N, which is convex. It is worth noting that if the point on $\mathcal{C}$ is in the image of multiple top dimensional simplices of N, the preimage of the point remains convex because the maps of these simplices agree on their common faces. Hence f is surjective and the pre-image of any point is contractible. Maps with these properties are called simple and are always homotopy equivalences [2].

□

2.2 Proximity

We start with the following lemma:

Lemma 2.2. *Consider points x, q, p in $\mathbb{R}^n$ such that $d(x,p) < d(x,q)$. Let $y = x + s * \vec{n}$ where $\vec{n}$ is a unit vector such that $(q-p)\cdot\vec{n} > 0$ ($\cdot$ is the dot vector) and*

$$s = \frac{d(x,q)^2 - d(x,p)^2}{2(q-p)\cdot\vec{n}} \tag{7}$$

See Figure 2. Then $d(y,p) = d(y,q)$.

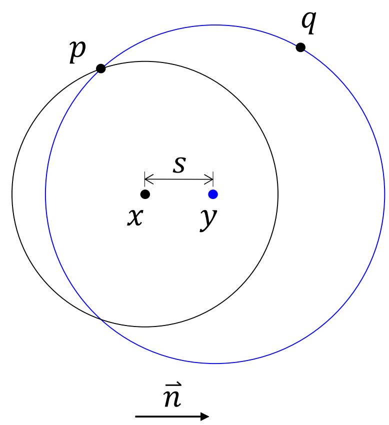

Figure 2: Notations for Lemma 2.2.

Proof. Noting that $d(\vec{u})^2 = \vec{u}\cdot\vec{u}$ for any vector $\vec{u}$, we can derive that

$$\begin{aligned}
d(y,p)^2 &= (x - p + s*\vec{n})\cdot(x - p + s*\vec{n}) && \text{(by Definition of } y) && (8)\\
&= d(x,p)^2 + 2s*(x-p)\cdot\vec{n} + s^2 && && (9)\\
&= d(x,p)^2 + 2s*(q-p)\cdot\vec{n} + 2s*(x-q)\cdot\vec{n} + s^2 && && (10)\\
&= d(x,p)^2 + d(x,q)^2 - d(x,p)^2 + 2s*(x-q)\cdot\vec{n} + s^2 && \text{(by Definition of } s) && (11)\\
&= d(x,q)^2 + 2s*(x-q)\cdot\vec{n} + s^2 && && (12)\\
&= (x - q + s*\vec{n})\cdot(x - q + s*\vec{n}) && && (13)\\
&= d(y,q)^2 && \text{(by Definition of } y) && (14)
\end{aligned}$$

□

Consider a voxel shape $\mathcal{O}$ with voxel size h, boundary elements B and boundary vertices P. We denote by $\Gamma_B(x), \Gamma_P(x)$ the points on B and P nearest to x, respectively. Recall that the medial axis $\mathcal{M}$ consists of all points $x \in \mathcal{O}$ such that $|\Gamma_B(x)| \geq 2$. Based on Corollary 3.8 in the paper, the proof will use the following characterization of the voxel core $\mathcal{C}$: it consists of all points $x \in \mathcal{O}$ such that $|\Gamma_P(x)| \geq 2$ and $\Gamma_P(x)$ do not all lie on the same element (edge or face) of B.

We shall prove the Hausdorff distance bound in each direction, first from $\mathcal{C}$ to $\mathcal{M}$ (Theorem 2.3) and then from $\mathcal{M}$ to $\mathcal{C}$ (Theorem 2.4).

Theorem 2.3. *For any $x \in \mathcal{C}$, $d(x, \mathcal{M}) \leq \frac{1}{4}h$*

Proof. If $\Gamma_P(x) \subseteq \Gamma_B(x)$, and since $x \in \mathcal{O}$, it follows that $x \in \mathcal{M}$, and hence we have the trivial bound of $d(x,\mathcal{M}) = 0$. Otherwise, let $e \in B$ and point $p \in e$ such that $d(x,p) = d(x,e) = d(x,B)$. That is, p is the closes point on the boundary to x, and e is its containing element. Since $\Gamma_P(x)$ do not all lie on the same element, there is some $q \in \Gamma_P(x)$ that is not a vertex of e. Note that $d(x,v) \geq d(x,q) > d(x,p)$ (the first inequality holds because q is the closest to x among all boundary vertices, and the second inequality holds because p is the closest point to x on the boundary surface). See Figure 3 (a) for an illustration.

We will separately consider the cases where e is either a face or an edge.

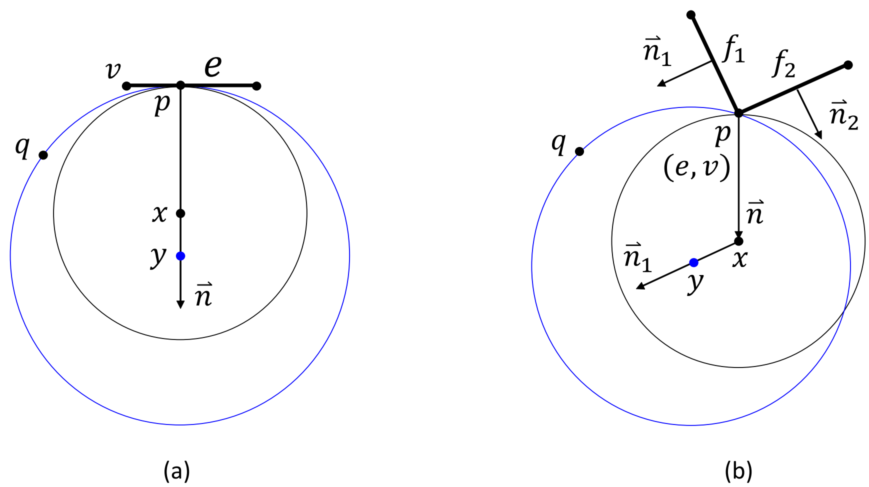

Figure 3: Notations used in the proof for face case (a) and edge case (b). In (b), the view is straight onto the edge e (hence e coincides with point p and vertex v).

Case 1: e is a face. We will first show that points x, q are on the same side of the supporting plane of e. Let v be the vertex of e closest to p. We have

$$d(p,v)^2 = d(x,v)^2 - d(x,p)^2 \geq d(x,q)^2 - d(x,p)^2 \tag{15}$$

Since $p \in e$, v is its closes vertex on e, and q is not a vertex of e, by Lemma 1.1, we have

$$d(p,q)^2 > d(p,v)^2 \tag{16}$$

Combining inequalities in (15, 16) we have

$$d(p,q)^2 + d(x,p)^2 > d(x,q)^2 \tag{17}$$

which implies that $(q-p)\cdot(x-p) > 0$. That is, both x, q lie "under" the supporting plane of e (see Figure 3 (a)).

Let $\vec{n}$ be the normalized vector of $(x-p)$, which is orthogonal to e and points inside $\mathcal{O}$. By Lemma 2.2, there exists some point y along the ray shot from x in the direction $\vec{n}$ such that $d(y,p) = d(y,q)$ and

$$d(x,y) = \frac{d(x,q)^2 - d(x,p)^2}{2(q-p)\cdot\vec{n}} \tag{18}$$

See Figure 3 (a). We can derive an upper bound of $d(x,y)$ as follows. Since e is a square of side length h, we know $d(p,v) \leq (\sqrt{2}/2)h$. Together with Equation 15, we bound the numerator by

$$d(x,q)^2 - d(x,p)^2 \leq h^2/2. \tag{19}$$

On the other hand, while we have shown that $(q-p)\cdot\vec{n} > 0$ is positive, we further make a crucial observation that, as both q and e are elements of the voxel grid, q is at least h distance away from the supporting plane of e. That is, we can bound the denominator by

$$(q-p)\cdot\vec{n} \geq h. \tag{20}$$

Given these two bounds, we have

$$d(x,y) \leq \frac{1}{4}h. \tag{21}$$

We next show that the closed segment $[x,y]$ must contain some point of the medial axis $\mathcal{M}$. Suppose the half-open segment $[x,y)$ contains no point of $\mathcal{M}$, we will show that $y \in \mathcal{M}$. Since p is the closest point of x

on the boundary B and $[x, y]$ follows the normal vector at p, and since the segment $[x, y)$ avoids the medial axis, as we move along the segment from x to y, we stay inside $\mathcal{O}$ and p remains the closest point on the boundary. Hence $y \in \mathcal{O}$ and $d(y, q) = d(y, p) = d(y, B)$, implying $y \in \mathcal{M}$. This leads to the bound

$$d(x, \mathcal{M}) \leq d(x, y) \leq \frac{1}{4} h \tag{22}$$

Case 2: e is an edge. Let the faces sharing e be f_1, f_2. We will first show that points x, q are on the same side of the supporting plane of at least one of these faces. Denote the inward unit normal vectors of the two faces as $\vec{n}_1$ and $\vec{n}_2$, and $\vec{n}$ as the normalized vector of $(x - p)$ (see Figure 3 (b)). As p is x's nearest point on the boundary, $\vec{n}$ is a convex combination of $\vec{n}_1$ and $\vec{n}_2$. That is,

$$\vec{n} = a_1 \vec{n}_1 + a_2 \vec{n}_2, \tag{23}$$

where a_1, a_2 are strictly positive. On the other hand, using the same argument in Case 1 (Equations 15,16,17), we can show that $(q - p) \cdot \vec{n} > 0$. Substituting into $\vec{n}$, we have

$$a_1 (q - p) \cdot \vec{n}_1 + a_2 (q - p) \cdot \vec{n}_2 > 0 \tag{24}$$

Since both a_1, a_2 are positive, at least one of the two quantities, $(q-p) \cdot \vec{n}_1$ or $(q-p) \cdot \vec{n}_2$ needs to be positive to make the inequality hold. Let us assume that $(q-p) \cdot \vec{n}_1 > 0$, that is, both x, q lie "under" the supporting plane of face f_1.

By Lemma 2.2, there exists some point y along the ray shot from x in the direction $\vec{n_1}$ such that $d(y, p) = d(y, q)$ and

$$d(x, y) = \frac{d(x, q)^2 - d(x, p)^2}{2(q - p) \cdot \vec{n_1}} \tag{25}$$

See Figure 3 (b). We can derive an upper bound of $d(x, y)$ similarly as in Case 1. Let v be the vertex of e closest to v. Since e is an edge of length h, we know $d(p, v) \leq h/2$. We bound the numerator by

$$d(x, q)^2 - d(x, p)^2 \leq h^2/4. \tag{26}$$

On the other hand, as both q and f_1 are elements of the voxel grid, q is at least h distance away from the supporting plane of f_1. That is, we can bound the denominator by

$$(q - p) \cdot \vec{n_1} \geq h. \tag{27}$$

Given these two bounds, we have

$$d(x, y) \leq \frac{1}{8} h. \tag{28}$$

Following a similar argument as in Case 1, the closed segment $[x, y]$ must contain some point of the medial axis $\mathcal{M}$. Note that while $\vec{n}_1$ is not the normal vector at the closest point p, moving from x to y (assuming we avoid $\mathcal{M}$) stays in the region of $\mathcal{O}$ whose closest point is p. This leads to the bound

$$d(x, \mathcal{M}) \leq d(x, y) \leq \frac{1}{8} h \tag{29}$$

□

Theorem 2.4. *For any $x \in \mathcal{M}$, $d(x, \mathcal{C}) \leq \frac{\sqrt{3}}{2} h$*

Proof. If $\Gamma_B(x)$ contains two or more vertices of P, then it is easy to see that x is also on the voxel core $\mathcal{C}$ since $x \in \mathcal{O}$, $\Gamma_P(x)$ has at least two vertices (i.e., $\Gamma_B(x) \cap P$), and these vertices do not lie on the same boundary edge or face on B (otherwise that edge or face would be closer to x than its vertices). Hence we have the trivial bound of $d(x, \mathcal{C}) = 0$.

Otherwise, let E denote the set of boundary edges and faces that contain some point of $\Gamma_B(x)$ in their interior. Note that $E \neq \emptyset$ (otherwise $\Gamma_B(x)$ would contain two or more vertices of P). We first consider a special case when the elements of E are sharing common vertices that are also the closest boundary vertices to x (Case 1). Then we will consider the remaining situations (Case 2).

Case 1: Suppose E has two elements e_1, e_2 that share a common vertex $q \in \Gamma_P(x)$ (that is, q is the closest boundary vertex to x). We will prove $d(x,\mathcal{C}) \le \frac{\sqrt{3}}{2}h$ for each possible composition of e_1, e_2 as edges or faces:

Case 1.1: Both e_1, e_2 are faces. Since both faces are tangent to the maximal ball centered at x and they share at least one common vertex, e_1, e_2 must share a common edge and bound a common voxel (Figure 4 (a)). Since the orthogonal projection of x onto e_1 (or e_2) must lie completely within e_1 (or e_2), x is contained within that voxel. Observe that the center of the voxel, c, is on the vertex core $\mathcal{C}$. This is because c has at least 6 nearest vertices on P (e.g., vertices of e_1, e_2), these vertices do not all lie on the same boundary element, and $c \in \mathcal{O}$ (because the voxel is either completely inside or outside $\mathcal{O}$, and $x \in \mathcal{O}$ is in the voxel). Hence $d(x,\mathcal{C}) \le d(x,c) \le \frac{\sqrt{3}}{2}h$.

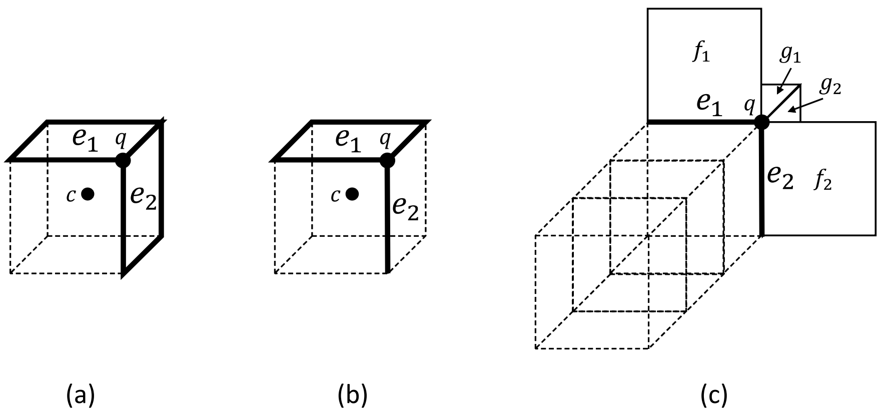

Figure 4: Illustrations for Cases 1.1 (a), 1.2 (b), and 1.3 (c).

Case 1.2: e_1 is a face and e_2 is an edge. Using a similar argument as above, both e_1, e_2 must bound a common voxel that contains x (Figure 4 (b)), and the center c of that voxel lies on $\mathcal{C}$. Again, we obtain $d(x,\mathcal{C}) \le d(x,c) \le \frac{\sqrt{3}}{2}h$.

Case 1.3: Both e_1, e_2 are edges. Since e_1 (and similarly for e_2) is tangent to the maximal ball centered at x, e_1 is shared by two boundary faces (f_1, g_1), and x is confined to a "slab" of voxels that lie under both supporting planes of f_1, g_1 and sandwiched between the two planes, each containing one vertex of e_1, that are orthogonal to e_1. Since e_1, e_2 share a common vertex, the only possible configuration is depicted in Figure 4 (c), where the boundary faces f_1, g_1 containing e_1 and the boundary faces f_2, g_2 containing e_2 are distinct, and two of them (e.g., g_1, g_2) share a common edge. Let L be the plane containing both e_1, e_2. x is then confined to a "string" of voxels that lie on the opposite side of L as g_1, g_2, such that orthogonal projection of these voxels onto L is the voxel face containing both e_1, e_2 (this string of voxels are shown in dotted lines in Figure 4 (c)).

First, consider the situation where x lies in the voxel containing e_1, e_2 (i.e., the first voxel in the string). Observe that the center c of the voxel is on $\mathcal{C}$, since it has at least three nearest vertices in P (e.g., those vertices of e_1, e_2) that do not lie on the same boundary element (since the voxel face containing e_1, e_2 is not a boundary face), and c is in $\mathcal{O}$ (because $x \in \mathcal{O}$ is in that voxel). Hence we obtain a similar bound as in the previous cases: $d(x,\mathcal{C}) \le d(x,c) \le \frac{\sqrt{3}}{2}h$.

Next, suppose x lies in some other voxel in the string. If the nearest vertex in P to the center c of that voxel is q, we can argue similarly as above that $c \in \mathcal{C}$ and $d(x,\mathcal{C}) \le d(x,c) \le \frac{\sqrt{3}}{2}h$. Otherwise, suppose q is not the nearest boundary vertex to c. Since q is the nearest boundary vertex to x, the (open) line segment (x, c) must intersect some face f of the Voronoi cell of q in the Voronoi diagram $VD(P)$. Note that f cannot be dual to a boundary edge. This is because the only boundary edges containing q whose dual Voronoi faces intersect this string of voxels are e_1, e_2. However, the supporting planes of dual Voronoi faces of e_1, e_2 (which bisect the two vertices of each edge) do not intersect with the line segment (x, c). In addition, since (x, c) is confined within the voxel of c, which lies in $\mathcal{O}$, f has at least one point in $\mathcal{O}$. By definition of the voxel core, f must be a face in $\mathcal{C}$. Let the intersection between (x, c) and f be y. Since $y \in \mathcal{C}$, we obtain $d(x, \mathcal{C}) \leq d(x, y) < d(x, c) \leq \frac{\sqrt{3}}{2}h$.

Case 2: Suppose the situation of Case 1 does not arise. That is, there does not exist two elements of E that share a vertex in $\Gamma_P(x)$. Then there are two possibilities:

- E contains a single element e (edge or face). Since $\Gamma_B(x)$ has at least two points, and one of them is on e, $\Gamma_B(x)$ contains at least one boundary vertex $q \in P$, and hence $q \in \Gamma_P(x)$. Note that, since the point on e closest to x is interior to e, and x is equal-distant to e and to q, q cannot be a vertex of e.

- E contains multiple elements, and they do not mutually share any vertex in $\Gamma_P(x)$. Therefore, for any $q \in \Gamma_P(x)$, there exists at least one element $e \in E$ such that q is not a vertex of e.

In either case, we can find a pair $\{e, q\}$ where e is a boundary edge or boundary face that contains a closest point to x, q is the closest boundary vertex to q, and q is not a vertex of e. Let p be the interior point of e that is closest to x, and v the vertex of e that is closest to p. Note that $d(x, v) \geq d(x, q) \geq d(x, p)$ (the first inequality holds because q is the closest to x among all boundary vertices P, and the second inequality holds because p is the closest point to x on the entire boundary surface B). See Figure 5 (a).

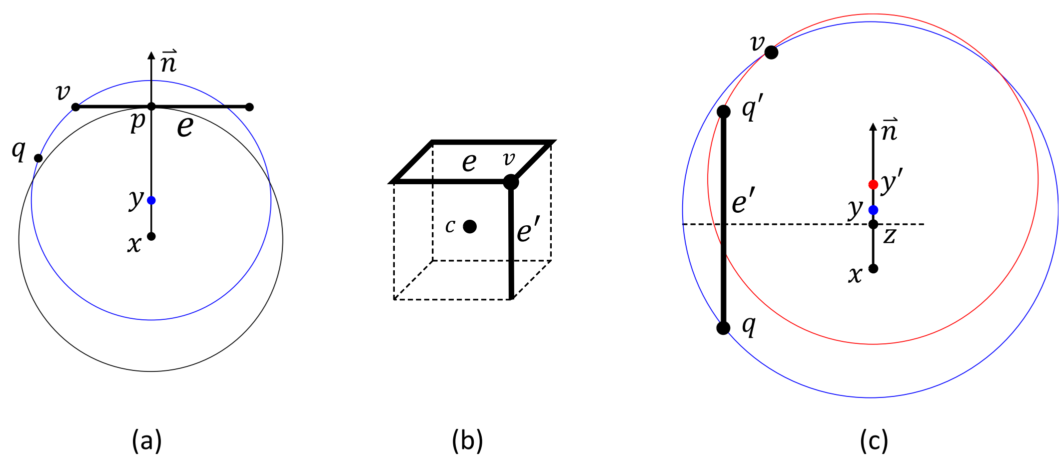

Figure 5: Illustrations for Cases 2.1 (a), 2.1.1 (b), and 2.1.2 (c).

We will separately consider the cases where e is either a face or an edge. When e is a face (Case 2.1), we prove that $d(x, \mathcal{C}) \leq \frac{\sqrt{3}}{2}h$. When e is an edge (Case 2.2), we obtain $d(x, \mathcal{C}) \leq h/8$.

<u>Case 2.1: e is a face.</u> Using the same argument as in Case 1 of the proof of Theorem 2.3 (Equations 15,16,17), we can show that both x and vertex q lie "under" the supporting plane of e (i.e., on the inward side of e; see Figure 5 (a)). We will first consider the situation where vertex v is contained in some boundary edge under the supporting plane of e as x (Case 2.1.1), and then examine the situation otherwise (Case 2.1.2).

Case 2.1.1: Suppose there is a boundary edge e' containing vertex v, such that the other vertex of e' lies under the supporting plane of e (see Figure 5 (b)). In this configuration, x is confined to the voxel bounded by e, e'. This is because x's closest point on e (i.e., p) is restricted to a quadrant of e containing v, and x is no further from e than from e'. Following the same argument as in Cases 1.1 and 1.2, the center point c of the voxel is on $\mathcal{C}$, and hence $d(x, \mathcal{C}) \leq d(x, c) \leq \frac{\sqrt{3}}{2}h$.

Case 2.1.2: Suppose the vertex v is not contained in any boundary edge under the supporting plane of e. Let $\vec{n}$ be the unit outward normal vector of face e. Our goal is to show that the ray shot from x in the direction of $\vec{n}$ will intersect with $\mathcal{C}$ at a location not too far from x. In fact, consider the point y on that ray such that $d(y, q) = d(y, v)$ (see Figure 5 (a)). By Lemma 2.2,

$$d(x, y) = \frac{d(x, v)^2 - d(x, q)^2}{2(v - q) \cdot \vec{n}} \tag{30}$$

Since $d(x,q) \geq d(x,p)$, $d(v,p)^2 = d(x,v)^2 - d(x,p)^2$, and $d(v,p)^2 \leq h^2/2$ (since v is the vertex of the square-shaped face e closest to p), the numerator of the right-hand side of Equation 30 is upper bounded by $h^2/2$. On the other hand, since q lies under e and both q, e are voxel elements, q is at a distance h or more from the supporting plane of e, which yields $2(v-q)\cdot\vec{n} \geq 2h$. As a result,

$$d(x,y) \leq \frac{1}{4}h \tag{31}$$

We next show that the line segment $[x,y]$ must intersect $\mathcal{C}$ under the following condition (denoted by $\diamondsuit$ for convenience): *Let $e' = \{q,q'\}$ be the voxel edge containing q in the direction of $\vec{n}$, then either e' is not a boundary edge, or the dual Voronoi face of e' in $VD(P)$ does not intersect the segment* $[x,y]$. Figure 5 (c) illustrates a violation of the condition, where segment $[x,y]$ intersects with the dual Voronoi face of e'. Suppose condition $\diamondsuit$ is met, to prove that $[x,y]$ intersects with $\mathcal{C}$, it suffices to show that, if the half-open segment $[x,y)$ avoids $\mathcal{C}$, then $y \in \mathcal{C}$. To do so, we need to show that $y \in \mathcal{O}$, $\{q,v\} \subseteq \Gamma_P(y)$, and q,v do not lie on the same boundary element. We prove each of these statements below:

1. Suppose $y \notin \mathcal{O}$. Hence y must lie outside the maximal ball centered at x with radius $d(x,p)$, implying $d(x,p) < d(x,y) \leq h/4$. Therefore x is at most $h/4$ away from the supporting plane of e, meaning that x is located within the voxel that is directly under face e and x is closer to e than to the bottom face of the voxel. This makes vertex v closer to x than all vertices on the voxel grid (including q), which contradicts to the fact that $q \in \Gamma_P(x)$.

2. Since $d(y,q) = d(y,v)$, we only need to show that $q \in \Gamma_P(y)$. Suppose $q \notin \Gamma_P(y)$. Since $q \in \Gamma_P(x)$, the (open) line segment (x,y) must intersect a face f of the Voronoi cell of q in $VD(P)$. Since $[x,y)$ avoids $\mathcal{C}$, f must be dual to a boundary edge that contains q. Among the six voxel edges containing q, the only edge whose dual Voronoi face can possibly intersect the line segment (x,y) is e' (which is parallel to the segment). However, such possibility is excluded by the condition $\diamondsuit$.

3. Suppose q,v lie on the same boundary element. Since q is under the supporting plane of e, this common boundary element must contain a boundary edge containing v that is under e, which violates the assumption of Case 2.1.2.

The argument above shows that $[x,y]$ contains some point of $\mathcal{C}$, and hence $d(x,\mathcal{C}) \leq d(x,y) \leq h/4$.

Next we consider the situation when condition $\diamondsuit$ is violated. That is, voxel edge $e' = \{q,q'\}$ is a boundary edge, and the dual Voronoi face of e' intersects with the line segment $[x,y]$, say at point z. Consider the point y' on the ray shot from x in the direction of $\vec{n}$ such that $d(y',q') = d(y',v)$ (see Figure 5 (c)). By Lemma 2.2,

$$d(x,y') = \frac{d(x,v)^2 - d(x,q')^2}{2(v-q')\cdot\vec{n}} \tag{32}$$

Since $d(x,q') \geq d(x,q)$, we can bound of the numerator of the right-hand side of Equation 32 as $d(x,v)^2 - d(x,q')^2 \leq d(x,v)^2 - d(x,q)^2 \leq h^2/2$. To bound the denominator, we need to show that q' is not on the supporting plane of e (which would imply that q' is at distance of h or more under the plane of e, and hence $2(v-q')\cdot\vec{n} \geq 2h$). Suppose that q' is on the supporting plane of e. This implies $d(z,p) = h/2$ (because z lies on the bisecting plane between q,q', and q is under e). Since $d(x,y) \leq h/4$ and z is on the line segment $[x,y]$, we have $d(x,p) \leq 3h/4$. Hence x lies in the voxel directly under e. Since q is the closest boundary vertex to x, q' must coincide with v, which violates the assumption of Case 2.1.2. As a result,

$$d(x,y') \leq \frac{1}{4}h \tag{33}$$

We show that the line segment $[z,y']$ must intersect with $\mathcal{C}$. It suffices to show that, if the half-open segment $[z,y')$ avoids $\mathcal{C}$, then $y' \in \mathcal{C}$. In fact, we can use the same arguments as those used under condition $\diamondsuit$ (simply replacing y,q by y',q' in the three bullets) to show that $y' \in \mathcal{O}$ and q',v do not lie on a same boundary element. To show $q' \in \Gamma_P(y')$ (and therefore $v \in \Gamma_P(y')$), we observe that y' is at most $h/4$ "above" the bisecting plane of q,q' (refer to Figure 5 (c)). Consider the other voxel edge $e'' = \{q',q''\}$ that contains q' and runs parallel to $\vec{n}$. The segment $[z,y')$ lies strictly between the bisecting planes of q,q' and of q',q'', and hence the segment avoids the dual Voronoi face of e'' in the case that e'' is a boundary edge.

Since $q' \in \Gamma_P(z)$ (because z is on a face bounding the Voronoi cell of q'), and because segment $[z, y')$ avoids all faces of the Voronoi cell of q', we conclude that $q' \in \Gamma_P(y')$. Now that the segment $[z, y']$ contains some point of $\mathcal{C}$, and the segment $[x, y']$ contains $[z, y']$, we obtain $d(x, \mathcal{C}) \leq d(x, y') \leq h/4$.

Case 2.2: e is an edge. Denote the two boundary faces sharing e as f_1, f_2 and their unit outward normals as $\vec{n}_1, \vec{n}_2$. Using the same argument as in Case 2 of the proof of Theorem 2.3 (Equations 23,24), we can show that both x and vertex q lie "under" the supporting plane of at least one boundary face, say f_1 (see Figure 6 (a)). This implies that q is at a distance of h or more under the supporting plane of f_1. We will consider the cases where vertices q, v lie on some common boundary element (Case 2.2.1) or not (Case 2.2.2).

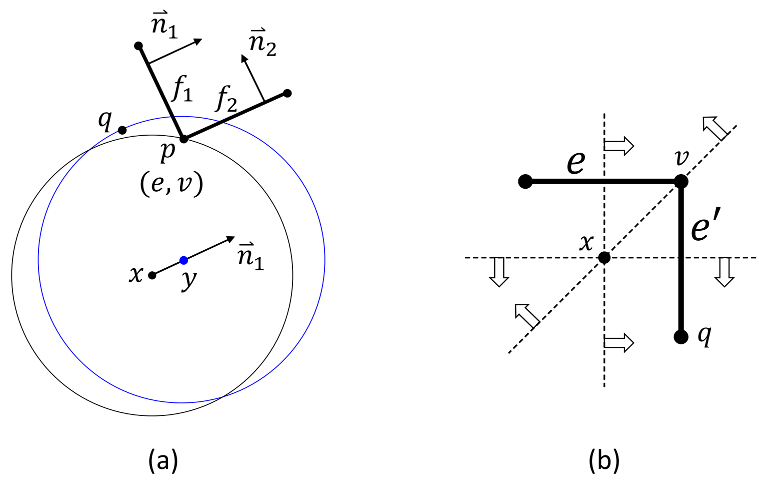

Figure 6: Illustrations for Cases 2.2 (a) and 2.2.1 (b). In (a), the view is straight onto the edge e (hence e coincides with point p and vertex v).

Case 2.1.1: Suppose $e' = \{q, v\}$ is a boundary edge. Since e is closer to x than e', x lies in the half-space (containing e) defined by the bisecting plane of e, e'. Since v is the vertex on e closer to x, x lies in the half-space (containing v) defined by the bisecting plane of the two vertices of e. Finally, since q is the closest boundary vertex to x, x lies in the half-space (containing q) defined by the bisecting plane between v, q. The intersection of these three half-spaces is the line that is the loci of points equidistant to v, q and the other vertex of e (see illustration in Figure 6 (b)). Hence x has at least three nearest vertices in P, which do not all lie on the same boundary face. Since $x \in \mathcal{M} \subset \mathcal{O}$, we have $x \in \mathcal{C}$. This leads to the trivial bound $d(x, \mathcal{C}) = 0$.

It is not possible for q, v to lie on a common boundary face f but do not share a common edge. If so, x would lie in the intersection of three half-spaces defined by the bisecting cone (or plane) between e, f, the bisecting plane between v and the other vertex of e, and the bisecting plane between v, q. It can be verified that such intersection is empty for all possible configurations of q and f (there are 4 configurations in total, since f must contain the voxel edge that contains v and lies under the plane of f_1).

Case 2.1.2: Now suppose q and v are not vertices of a common boundary element. Our goal is to show that the ray shot from x in the direction of $\vec{n}_1$ will intersect with $\mathcal{C}$ at a location not too far from x. Consider the point y on that ray such that $d(y, q) = d(y, v)$ (see Figure 6 (a)). By Lemma 2.2,

$$d(x, y) = \frac{d(x, v)^2 - d(x, q)^2}{2(v - q) \cdot \vec{n}_1} \tag{34}$$

Following the same argument in Case 2.1.2 (below Equation 30), and note that $d(v, p)^2 \leq h^2/4$ (since v is the vertex of edge e closer to p), we have

$$d(x, y) \leq \frac{1}{8} h \tag{35}$$

We show that the line segment $[x, y]$ intersects $\mathcal{C}$. It suffices to show that, if the half-open segment $[x, y)$ avoids $\mathcal{C}$, then $y \in \mathcal{C}$. To this end, we need to show that $y \in \mathcal{O}$, $\{q, v\} \subseteq \Gamma_P(y)$, and q, v do not lie on the same boundary element. Since the last statement is the assumption of Case 2.1.2, we only need to prove the first two statements:

1. Suppose $y \notin \mathcal{O}$. Hence y lies outside the maximal ball centered at x with radius $d(x, p)$, implying $d(x, p) < d(x, y) \leq h/8$. Therefore x is at most $h/8$ away from the edge e, which makes v closer to x than all vertices on the voxel grid (including q), thus contradicting $q \in \Gamma_P(x)$.

2. Since $d(y, q) = d(y, v)$, it suffices to show that $q \in \Gamma_P(y)$. Suppose $q \notin \Gamma_P(y)$. Since $q \in \Gamma_P(x)$, the (open) line segment (x, y) must intersect a face of the Voronoi cell of q in $VD(P)$. Since $[x, y)$ avoids $\mathcal{C}$, the Voronoi face must be dual to the boundary edge $e' = \{q, q'\}$ that is parallel to $[x, y)$. Let the intersection of the dual Voronoi face of e' and segment $[x, y)$ be z (refer to Figure 5 (c)). Note that $d(x, z) < d(x, y)$. To derive a contradiction, we observe that

$$d(x,q)^2 = d(x,e')^2 + (\frac{h}{2} - d(x,z))^2 \geq d(x,p)^2 + (\frac{h}{2} - d(x,z))^2 \tag{36}$$

Substituting the above into Equation 34, we have

$$\begin{aligned} d(x,y) &\leq \frac{d(x,v)^2 - d(x,p)^2 - (\frac{h}{2} - d(x,z))^2}{2h} && (37)\\ &\leq \frac{\frac{h^2}{4} - (\frac{h}{2} - d(x,z))^2}{2h} && (38)\\ &= \frac{h\ d(x,z) - d(x,z)^2}{2h} && (39)\\ &\leq \frac{d(x,z)}{2} && (40)\\ & && (41) \end{aligned}$$

which leads to the contradiction that $d(x, z) < d(x, y) \leq (d(x, z))/2$.

The argument above shows that $[x, y]$ contains some point of $\mathcal{C}$, and hence $d(x, \mathcal{C}) \leq d(x, y) \leq h/8$.

□

The combination of Theorems 2.3 and 2.4 gives the Hausdorff distance bound of $\frac{\sqrt{3}}{2}h$ between $\mathcal{C}$ and $\mathcal{M}$ in Theorem 3.5 of the paper.

3 Voxelization of smooth surfaces

We give proofs for Theorems 3.9 and 3.10 in the paper. We consider an open set $\mathcal{O} \in \mathbb{R}^3$ bounded by a C^2 continuous manifold surface B with a positive reach r, and a voxelization $\mathcal{O}_h$ with voxel size h is the interior of the union of all voxels whose centers lie in $\mathcal{O}$. Denote by B_h, P_h the boundary elements and boundary vertices of $\mathcal{O}_h$.

3.1 Proximity

We first recall the result due to Lachaud and Thibert [3] that bounds the Hausdorff distance between the smooth boundary B and the voxelized boundary B_h:

Theorem 3.1 (Theorem 1 in [3])**.** *For any* $h < \frac{2\sqrt{3}}{3}r$, $d_H(B_h, B) \leq \frac{\sqrt{3}}{2}h$.

Building on this distance bound, we next give the distance bounds between two pairs of structures: between the voxel boundary vertices P_h and the smooth boundary B, and between the voxel shape O_h and the open set O.

Theorem 3.2 (Theorem 3.9 in paper)**.** *For any* $h < \frac{2\sqrt{3}}{3}r$,

1. $d_H(P_h, B) \leq \frac{\sqrt{2}+\sqrt{3}}{2}h$, *and*

2. $d_H(\mathcal{O}_h, \mathcal{O}) \leq \frac{\sqrt{3}}{2}h$.

Proof. We prove each statement separately.

1. Consider a voxel boundary vertex $p \in P_h$. Since $p \in B_h$, by Theorem 3.1, we immediately have $d(p, B) \leq (\sqrt{3}/2)h$. Next, consider any point $x \in B$. By Theorem 3.1, there exists a point $y \in B_h$ such that $d(x, y) \leq (\sqrt{3}/2)h$. Let e be the voxel element (face, edge or vertex) of B_h that contains y. Then there exists some vertex p of e such that $d(y, p) \leq (\sqrt{2}/2)h$. Since $p \in P_h$, we have

$$d(x, P_h) \leq d(x, p) \leq d(x, y) + d(y, p) \leq \frac{\sqrt{2}+\sqrt{3}}{2}h.$$

2. Consider a point $x \in \mathcal{O}_h$. Let v be the voxel of $\mathcal{O}_h$ containing x. Since the center of the voxel, c, must lie in $\mathcal{O}$, we immediately have $d(x, \mathcal{O}) \leq d(x, c) \leq (\sqrt{3}/2)h$.

   Next, consider a point $x \in \mathcal{O}$. Let $y \in B$ be the closes point to x on the boundary B. We define a sphere S as follows. If $d(x, y) \leq (\sqrt{3}/2)h$, we define S as the ball with radius $(\sqrt{3}/2)h$ that is tangent to B at y such that the vector from y to the center o of S points to the inside of O (see Figure 7 left). Since $r > (\sqrt{3}/2)h$, and by Lemma 1 of [3], S lies completely inside $\mathcal{O}$ (except at y). If $d(x, y) > (\sqrt{3}/2)h$, define S as the ball centered at x with radius $(\sqrt{3}/2)h$ (see Figure 7 right). In either case, S has radius $(\sqrt{3}/2)h$, lies in $\mathcal{O}$ (except possibly for a single point on the boundary of S), and contains x. We next argue that the center o of S lies in the closure of $\mathcal{O}_h$. Consider the (closed) voxel v that contains o and let c be the centroid of v (see Figure 7). Since the c is no farther than $(\sqrt{3}/2)h$ away from o, c lies in S. Since S lies in $\mathcal{O}$, c also lies in $\mathcal{O}$ (if c coincides with the single point on the boundary of S that does not lie in $\mathcal{O}$, o must be a vertex of v, and we will replace v by any other voxel incident to o). Hence voxel v lies in the closure of $\mathcal{O}_h$, implying o lies in this closure too. Hence we have $d(x, \mathcal{O}_h) \leq d(x, o) \leq (\sqrt{3}/2)h$.

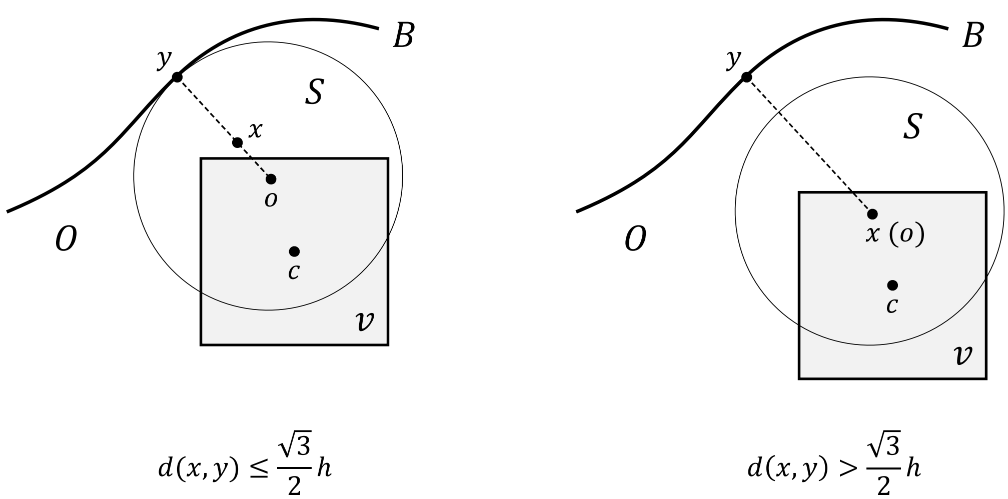

Figure 7: Notations for the proof of Theorem 3.2.

□

3.2 Homotopy equivalence

Our proof builds on prior results by Stelldinger et al. [5] concerning the topology equivalence between a smooth shape $\mathcal{O}$ and its reconstruction $\mathcal{O}'$ on a Cartesian grid. Note that the straightforward reconstruction

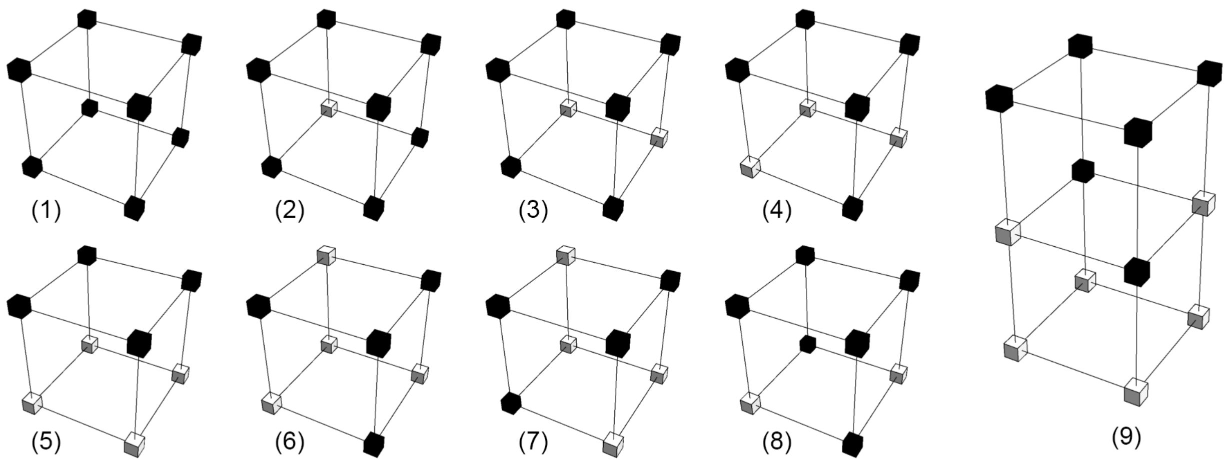

Figure 8: Configurations (1-8) of dual cubes represented by the inside/outside states of the cube corners (which are centroids of voxels). When the voxel resolution is sufficiently high ($h < \frac{2\sqrt{3}}{3}r$), these are the only possible configurations, and configuration (8) only appears in pairs as shown in (9).

as the union of voxels (i.e., the closure of voxel shape $\mathcal{O}_h$) may not be homeomorphic to $\mathcal{O}$ regardless of the sampling resolution, due to the presence of non-manifold edges and vertices on the voxel boundary B that do not vanish as voxels get smaller. However, Stelldinger et al. show that any reconstruction that satisfies certain local manifoldness and separation properties is homeomorphic to $\mathcal{O}$ under mild sampling conditions. To prove the homotopy equivalence between $\mathcal{O}_h$ and $\mathcal{O}$, we will show that $\mathcal{O}_h$ can be deformed in a homotopy-preserving way to a reconstruction $\mathcal{O}'$ that satisfies Stelldinger's properties.

We will first recall results from Stelldinger et al. We define the *dual* of the voxel grid as elements with complementary dimensionality on a grid that is offset by $h/2$ in each direction from the voxel grid. For example, the dual of a voxel vertex p, denoted by $\widetilde{p}$, is a cube formed by the centroids of the 8 voxels sharing p. A *configuration* of $\widetilde{p}$ refers to the inside/outside states of the 8 corners of $\widetilde{p}$ with respect to the shape $\mathcal{O}$. Two dual cubes are said to have the same configuration if they can be identify by rotation, reflection or switching of inside/outside states at all cube corners. A key observation made by Stelldinger et al. is that, for sufficiently small h, the only possible dual cube configurations are those listed in Figure 8 (1-8). Furthermore, Configuration 8 always appears in pairs with complementary inside/outside states, as shown in Figure 8 (9). This observation allows the authors to show the following:

Theorem 3.3 (Definition 15 and Theorem 16 in [5])**.** *A shape $\mathcal{O}'$ reconstructed from a voxel grid whose size satisfies $h < \frac{2\sqrt{3}}{3}r$ is homeomorphic to $\mathcal{O}$ if the following three conditions are met:*

1. *Any dual cube with Configuration 1 in Figure 8 lies completely inside (resp. outside) $\mathcal{O}'$ if all corners of the dual cube are inside (resp. outside) $\mathcal{O}$.*

2. *The intersection between any dual cube in Configurations 2 to 7 in Figure 8 and the boundary of $\mathcal{O}'$ is homeomorphic to a disk, and it divides the dual cube into two parts, each homeomorphic to a ball. The part of the dual cube inside (resp. outside) $\mathcal{O}'$ contains all corners of the dual cube that are inside (resp. outside) $\mathcal{O}$.*

3. *Consider a cuboid formed by the union of a pair of face-adjacent dual cubes, both with Configuration 8 in Figure 8 but with complementary inside/outside states, as shown in Figure 8 (9). The intersection*

between any such cuboid and the boundary of $\mathcal{O}'$ is homeomorphic to a disk, and it divides the cuboid into two parts, each homeomorphic to a ball. The part of the cuboid inside (resp. outside) $\mathcal{O}'$ contains all corners of the two composing dual cubes that are inside (resp. outside) $\mathcal{O}$.

We next prove our result:

Theorem 3.4 (Theorem 3.10 in the paper)**.** *For any $h < \frac{2\sqrt{3}}{3}r$, $\mathcal{O}_h$ is homotopy equivalent to $\mathcal{O}$.*

Proof. Observe that the voxel shape $\mathcal{O}_h$ meets the first two conditions in Theorem 3.3 (see Figure 9 top) but fails the third condition: the boundary B of $\mathcal{O}_h$ inside a dual cube of Configuration 8 contains a non-manifold edge, and hence the portion of B inside the union of a pair of dual cubes of Configuration 8 is not homeomorphic to a disk (see Figure 9 (a,b)). But we can construct a slightly deformed version of $\mathcal{O}_h$, noted as $\mathcal{O}'$, that meets all three conditions. For each dual cube of Configuration 8, consider a pair of tetrahedra, each formed by the voxel vertex c at the center of the cube and the centroids of three cube faces, such that both tetrahedra lie inside the closure of $\mathcal{O}_h$. The two tetrahedra share the non-manifold edge (see Figure 9 (a)). To obtain $\mathcal{O}'$, we "flatten" both tetrahedra by moving c along the non-manifold edge until it is identified with the other end of the non-manifold edge, which is the center of the cube face that has two inside corners and two outside corners (see Figure 9 (c)). Clearly, this flattening defines a deformation retract from the two tetrahedra to two triangles. Since the rest of $\mathcal{O}_h$ is unchanged, $\mathcal{O}'$ is a deformation retract of $\mathcal{O}_h$, and hence is homotopy equivalent to $\mathcal{O}_h$. On the other hand, the boundary of $\mathcal{O}'$ inside the union of the two dual cubes of Configuration 8 is homeomorphic to a disk (see Figure 9 (d)). By Theorem 3.3, $\mathcal{O}'$ is homeomorphic to $\mathcal{O}$, and hence $\mathcal{O}_h$ is homotopy equivalent to $\mathcal{O}$. □

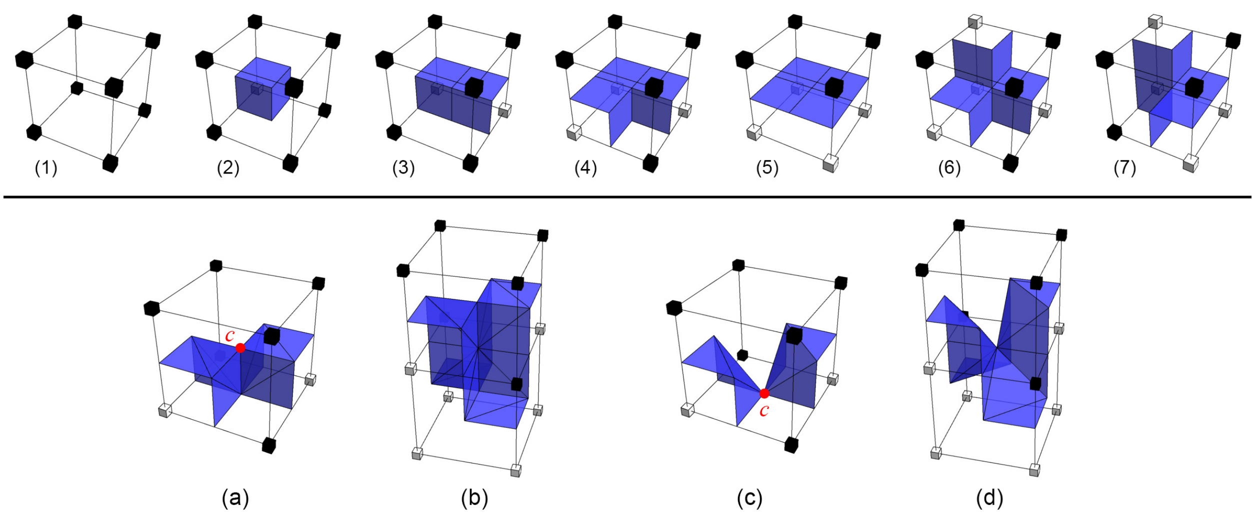

Figure 9: Top: boundary B of voxel shape $\mathcal{O}_h$ in a dual cube of Configurations 1-7 in Figure 8 is either empty or homeomorphic to a disk. Bottom: (a) The portion of B in a dual cube of Configuration 8 contains a non-manifold edge. (b) The portion of B in the cuboid formed by a pair of Configuration-8 dual cubes is not homeomorphic to a disk. (c) The boundary of the deformed shape $\mathcal{O}'$ has no non-manifold edge in the same dual cube. (d) The boundary of $\mathcal{O}'$ inside the cuboid is homeomorphic to a disk.

References

[1] Anders Björner, Bernhard Korte, and László Lovász. Homotopy properties of greedoids. *Advances in Applied Mathematics*, 6(4):447–494, 1985.

[2] Marshall M Cohen. Simplicial structures and transverse cellularity. *Annals of Mathematics*, pages 218–245, 1967.

[3] Jacques-Olivier Lachaud and Boris Thibert. Properties of gauss digitized shapes and digital surface integration. *Journal of Mathematical Imaging and Vision*, 54(2):162–180, 2016.

[4] Colin P Rourke and Brian Joseph Sanderson. *Introduction to piecewise-linear topology.* Springer Science & Business Media, 2012.

[5] Peer Stelldinger, Longin Jan Latecki, and Marcelo Siqueira. Topological equivalence between a 3d object and the reconstruction of its digital image. *IEEE transactions on pattern analysis and machine intelligence*, 29(1), 2007.