

Sampling Conditions for Conforming Voronoi Meshing by the VoroCrust Algorithm

Ahmed Abdelkader

University of Maryland, College Park MD, USA

akader@cs.umd.edu

Chandrajit L. Bajaj

University of Texas, Austin TX, USA

Mohamed S. Ebeida

Sandia National Laboratories, Albuquerque NM, USA

Ahmed H. Mahmoud

University of California, Davis CA, USA

Scott A. Mitchell

Sandia National Laboratories, Albuquerque NM, USA

John D. Owens

University of California, Davis CA, USA

Ahmad A. Rushdi

University of California, Davis CA, USA

Abstract

We study the problem of decomposing a volume bounded by a smooth surface into a collection of Voronoi cells. Unlike the dual problem of conforming Delaunay meshing, a principled solution to this problem for generic smooth surfaces remained elusive. VoroCrust leverages ideas from α -shapes and the power crust algorithm to produce unweighted Voronoi cells conforming to the surface, yielding the first provably-correct algorithm for this problem. Given an ϵ -sample on the bounding surface, with a weak σ -sparsity condition, we work with the balls of radius δ times the local feature size centered at each sample. The corners of this union of balls are the Voronoi sites, on both sides of the surface. The facets common to cells on opposite sides reconstruct the surface. For appropriate values of ϵ , σ and δ , we prove that the surface reconstruction is isotopic to the bounding surface. With the surface protected, the enclosed volume can be further decomposed into an isotopic volume mesh of fat Voronoi cells by generating a bounded number of sites in its interior. Compared to state-of-the-art methods based on clipping, VoroCrust cells are full Voronoi cells, with convexity and fatness guarantees. Compared to the power crust algorithm, VoroCrust cells are not filtered, are unweighted, and offer greater flexibility in meshing the enclosed volume by either structured grids or random samples.

2012 ACM Subject Classification Theory of computation → Randomness, geometry and discrete structures → Computational geometry

Keywords and phrases sampling conditions, surface reconstruction, polyhedral meshing, Voronoi

Digital Object Identifier 10.4230/LIPIcs.SoCG.2018.1

Related Version The full version is available online [1]. In addition, an accompanying multimedia contribution can be found in these proceedings [2].



© Ahmed Abdelkader, Chandrajit L. Bajaj, Mohamed S. Ebeida, Ahmed H. Mahmoud, Scott A. Mitchell, John D. Owens and Ahmad A. Rushdi;

licensed under Creative Commons License CC-BY

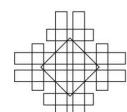
34th International Symposium on Computational Geometry (SoCG 2018).

Editors: Bettina Speckmann and Csaba D. Tóth; Article No. 1; pp. 1:1–1:16



Leibniz International Proceedings in Informatics

LIPIcs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



1 Introduction

Mesh generation is an important problem in computational geometry, geometric modeling, scientific computing and computer graphics. There has been a growing interest in polyhedral meshes as an alternative to tetrahedral or hex-dominant meshes. The main advantages of polyhedral meshes are higher degrees of freedom per element and fewer elements for the same number of vertices. This can be very useful in several numerical methods, e.g., finite volume [37], virtual element [17] and Petrov-Galerkin [39]. Within the class of polyhedral cells, Voronoi cells share several properties with tetrahedra, e.g., planar facets, convexity and positive Jacobians. In addition, the accuracy of a number of important solvers, e.g., the two-point flux approximation for conservation laws [37], greatly benefits from a conforming mesh which is *orthogonal* to its dual as naturally satisfied by Voronoi meshes. Such solvers play a crucial role in hydrology [47] and computational fluid dynamics [20].

VoroCrust is the first provably-correct algorithm for generating a volumetric Voronoi mesh whose boundary conforms to a smooth bounding surface, and with quality guarantees. A conforming volume mesh exhibits two desirable properties *simultaneously*: (1) a decomposition of the enclosed volume, and (2) a reconstruction of the bounding surface. Conforming Delaunay meshing is well-studied [26], but Voronoi meshing is less mature. A common practical approach to polyhedral meshing is to dualize a tetrahedral mesh and *clip*, i.e., intersect and truncate, each cell by the bounding surface [33, 40, 44, 48]. Unfortunately, clipping sacrifices the important properties of convexity and connectedness of cells, and requires costly constructive solid geometry operations. Restricting a Voronoi mesh to the surface before *filtering* its dual Delaunay facets is another approach [7, 31, 49], but filtering requires extra checks complicating its implementation and analysis; see also Figure 4. An intuitive approach is to locally mirror the Voronoi sites on either side of the surface [18, 32], but we are not aware of any robust algorithms with approximation guarantees in this category. In contrast to these approaches, VoroCrust is distinguished by its simplicity and robustness at producing true unweighted Voronoi cells, leveraging established libraries, e.g., Voro++ [46], without modification or special cases.

VoroCrust can be viewed as a principled mirroring technique, which shares a number of key features with the power crust algorithm [13]. The power crust literature [7, 8, 10, 12, 13] developed a rich theory for surface approximation, namely the ϵ -sampling paradigm. Recall that the power crust algorithm uses an ϵ -sample of unweighted points to place weighted sites, so-called *poles*, near the medial axis of the underlying surface. The surface reconstruction is the collection of facets separating power cells of poles on the inside and outside of the enclosed volume. Regarding samples and poles as primal-dual constructs, power crust performs a *primal-dual-dual-primal dance*. VoroCrust makes a similar dance where weights are introduced differently; the samples are weighted to define unweighted sites tightly hugging the surface, with the reconstruction arising from their unweighted Voronoi diagram. The key advantage is the freedom to place more sites within the enclosed volume without disrupting the surface reconstruction. This added freedom is essential to the generation of graded meshes; a primary virtue of the proposed algorithm. Another virtue of the algorithm is that all samples appear as vertices in the resulting mesh. While the power crust algorithm does not guarantee that, some variations do so by means of filtering, at the price of the reconstruction no longer being the boundary of power cells [7, 11, 30].

The main construction underlying VoroCrust is a suitable union of balls centered on the bounding surface, as studied in the context of non-uniform approximations [24]. Unions of balls enjoy a wealth of results [15, 22, 35], which enable a variety of algorithms [13, 21, 28].

Similar constructions have been proposed for meshing problems in the applied sciences with heuristic extensions to 3D settings; see [38] and the references therein for a recent example. Aichholzer et al. [6] adopt closely related ideas to construct a union of surface balls using power crust poles for sizing estimation. However, their goal was to produce a coarse homeomorphic surface reconstruction, and the connection to Voronoi meshing is absent. In contrast, VoroCrust aims at a decomposition of the enclosed volume into fat Voronoi cells conforming to an isotopic surface reconstruction with quality guarantees.

In a previous paper [4], we explored the related problem of generating a Voronoi mesh that conforms to restricted classes of piecewise-linear complexes, with more challenging inputs left for future work. The approach adopted in [4] does not use a union of balls and relies instead on similar ideas to those proposed for conforming Delaunay meshing [27, 42, 45]. Ultimately, we seek a conforming Voronoi mesher that can handle realistic inputs including a mix of smooth and sharp features, can estimate a sizing function and generate samples, and can guarantee the quality of the output mesh. This is the subject of a forthcoming paper [3] which describes the design and implementation of the complete VoroCrust algorithm.

In this paper, we present a theoretical analysis of an abstract version of the VoroCrust algorithm. This establishes the quality and approximation guarantees of its output for volumes bounded by smooth surfaces. A description of the algorithm we analyze is given next; see Figure 1 for an illustration in 2D, and also our accompanying multimedia contribution [2].

The abstract VoroCrust algorithm

1. Take as input a sample \mathcal{P} on the surface \mathcal{M} bounding the volume \mathcal{O} .
2. Define a ball B_i centered at each sample p_i , with a suitable radius r_i , and let $\mathcal{U} = \cup_i B_i$.
3. Initialize the set of sites \mathcal{S} with the corner points of $\partial\mathcal{U}$, \mathcal{S}^\uparrow and \mathcal{S}^\downarrow , on both sides of \mathcal{M} .
4. Optionally, generate additional sites $\mathcal{S}^{\downarrow\downarrow}$ in the interior of \mathcal{O} , and include $\mathcal{S}^{\downarrow\downarrow}$ into \mathcal{S} .
5. Compute the Voronoi diagram $\text{Vor}(\mathcal{S})$ and retain the cells with sites in $\mathcal{S}^\downarrow \cup \mathcal{S}^{\downarrow\downarrow}$ as the volume mesh $\hat{\mathcal{O}}$, where the facets between \mathcal{S}^\uparrow and \mathcal{S}^\downarrow yield a surface approximation $\hat{\mathcal{M}}$.

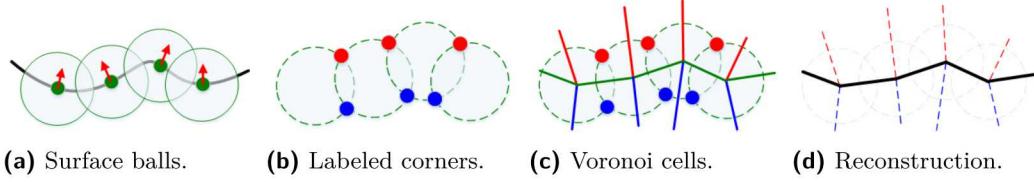


Figure 1 VoroCrust reconstruction, demonstrated on a planar curve.

In this paper, we assume \mathcal{P} is an ϵ -sample, with a weak σ -sparsity condition, and r_i is set to δ times the local feature size at p_i . For appropriate values of the parameters ϵ , σ and δ , we prove that $\hat{\mathcal{O}}$ and $\hat{\mathcal{M}}$ are isotopic to \mathcal{O} and \mathcal{M} , respectively. We also show that simple techniques for sampling within \mathcal{O} , e.g., octree refinement, guarantee an upper bound on the fatness of all cells in $\hat{\mathcal{O}}$, as well as the number of samples.

The rest of the paper is organized as follows. Section 2 introduces the key definitions and notation used throughout the paper. Section 3 describes the placement of Voronoi seeds and basic properties of our construction assuming the union of surface balls satisfies a structural property. Section 4 proves this property holds and establishes the desired approximation guarantees under certain conditions on the input sample. Section 5 considers the generation of interior samples and bounds the fatness of all cells in the output mesh. Section 6 concludes the paper. A number of proofs is deferred to the full version, available online [1]; see also the accompanying multimedia contribution in these proceedings [2].

2 Definitions and preliminaries

We assume the volume \mathcal{O} is a bounded open subset of \mathbb{R}^3 . The boundary of \mathcal{O} is a closed, bounded and smooth two-dimensional surface denoted by \mathcal{M} . The Euclidean distance between two points $p, q \in \mathbb{R}^3$ is denoted $\mathbf{d}(p, q)$. Throughout the paper, standard general position assumptions [36] are made to simplify the presentation. We proceed to recall the key definitions and notation used throughout the paper, following those in [13, 19, 24, 34, 35].

2.1 Sampling and approximation

We take as input a set of sample points $\mathcal{P} \subset \mathcal{M}$. A local scale or *sizing* is used to vary the sample density. Recall that the *medial axis* [13] of \mathcal{M} , denoted by \mathcal{A} , is the closure of the set of points in \mathbb{R}^3 with more than one closest point on \mathcal{M} . Hence, \mathcal{A} has one component inside \mathcal{O} and another outside. Each point of \mathcal{A} is the center of a *medial ball* tangent to \mathcal{M} at multiple points. Likewise, each point on \mathcal{M} has two tangent medial balls, not necessarily of the same size. With that, the *local feature size* at $x \in \mathcal{M}$ is defined as $\text{lfs}(x) = \mathbf{d}(x, \mathcal{A})$. The set \mathcal{P} is an ϵ -sample [9] if for all $x \in \mathcal{M}$ there exists $p \in \mathcal{P}$ such that $\mathbf{d}(x, p) \leq \epsilon \cdot \text{lfs}(x)$.

We desire an approximation of \mathcal{O} by a Voronoi mesh $\hat{\mathcal{O}}$, where the boundary $\hat{\mathcal{M}}$ of $\hat{\mathcal{O}}$ approximates \mathcal{M} . To define the type of approximations we desire, we recall a few definitions [24]. Two topological spaces are *homotopy-equivalent* if they have the same topology type. In other words, there is a one-to-one correspondence between their connected components, cycles, cavities, etc., as well as how these topological features are related. A stronger notion of topological equivalence is *homeomorphism*, which holds when there exists a continuous bijection with a continuous inverse from \mathcal{M} to $\hat{\mathcal{M}}$. Intuitively, two homeomorphic surfaces can be *smoothly deformed* into one another without tearing or self-intersection. The notion of isotopy better captures the topological equivalence for surfaces *embedded* in Euclidean space. Two surfaces $\mathcal{M}, \hat{\mathcal{M}} \subset \mathbb{R}^3$ are *isotopic* [16, 23] if there is a continuous mapping $F : \mathcal{M} \times [0, 1] \rightarrow \mathbb{R}^3$ such that for each $t \in [0, 1]$, $F(\cdot, t)$ is a homeomorphism from \mathcal{M} to $\hat{\mathcal{M}}$, where $F(\cdot, 0)$ is the identity of \mathcal{M} and $F(\mathcal{M}, 1) = \hat{\mathcal{M}}$. To capture the requirement that \mathcal{M} and $\hat{\mathcal{M}}$ should be close in terms of Euclidean distance, the notion of *Hausdorff distance* is defined as $d_H(X, Y) = \max\{\sup_{x \in X} \inf_{y \in Y} \mathbf{d}(x, y), \sup_{y \in Y} \inf_{x \in X} \mathbf{d}(x, y)\}$.

2.2 Diagrams and triangulations

The set of points defining a Voronoi diagram are traditionally referred to as *sites* or *seeds*. When approximating a manifold by a set of sample points of varying density, it is helpful to assign weights to the points reflective of their density. In particular, a point p_i with weight w_i , can be regarded as a ball B_i with center p_i and radius $r_i = \sqrt{w_i}$, i.e., $\mathbb{B}(p_i, r_i)$.

Recall that the *power distance* [35] between two points p_i, p_j with weights w_i, w_j is $\pi(p_i, p_j) = \mathbf{d}(p_i, p_j)^2 - w_i - w_j$. Unless otherwise noted, points are *unweighted*, having weight equal to zero. There is a natural geometric interpretation of the weight: all points q on the boundary of B_i have $\pi(p_i, q) = 0$, inside $\pi(p_i, q) < 0$ and outside $\pi(p_i, q) > 0$. Given a set of weighted points \mathcal{P} , this metric gives rise to a natural decomposition of \mathbb{R}^3 into the *power cells* $V_i = \{q \in \mathbb{R}^3 \mid \pi(p_i, q) \leq \pi(p_j, q) \forall p_j \in \mathcal{P}\}$. The *power diagram* $w\text{Vor}(\mathcal{P})$ is the cell complex defined by collection of cells V_i for all $p_i \in \mathcal{P}$.

The nerve [35] of a collection \mathcal{C} of sets is defined as $\mathcal{N}(\mathcal{C}) = \{X \subseteq \mathcal{C} \mid \cap T \neq \emptyset\}$. Observe that $\mathcal{N}(\mathcal{C})$ is an abstract simplicial complex because $X \in \mathcal{N}(\mathcal{C})$ and $Y \subseteq X$ imply $Y \in \mathcal{N}(\mathcal{C})$. With that, we obtain the *weighted Delaunay triangulation*, or *regular triangulation*, as $w\text{Del}(\mathcal{P}) = \mathcal{N}(w\text{Vor}(\mathcal{P}))$.

Alternatively, $\text{wDel}(\mathcal{P})$ can be defined directly as follows. A subset $T \subset \mathbb{R}^d$, with $d \leq 3$ and $|T| \leq d+1$ defines a d -simplex σ_T . Recall that the *orthocenter* [25] of σ_T , denoted by z_T , is the unique point $q \in \mathbb{R}^d$ such that $\pi(p_i, z_T) = \pi(p_j, z_T)$ for all $p_i, p_j \in T$; the *orthoradius* of σ_T is equal to $\pi(p, z_T)$ for any $p \in T$. The *Delaunay condition* defines $\text{wDel}(\mathcal{P})$ as the set of tetrahedra σ_T with an *empty orthosphere*, meaning $\pi(p_i, z_T) \leq \pi(p_j, z_T)$ for all $p_i \in T$ and $p_j \in \mathcal{P} \setminus T$, where $\text{wDel}(\mathcal{P})$ includes all faces of σ_T .

There is a natural duality between $\text{wDel}(\mathcal{P})$ and $\text{wVor}(\mathcal{P})$. For a tetrahedron σ_T , the definition of z_T immediately implies z_T is a *power vertex* in $\text{wVor}(\mathcal{P})$. Similarly, for each k -face f_S of $\sigma_T \in \text{wDel}(\mathcal{P})$ with $S \subseteq T$ and $k+1 = |S|$, there exists a dual $(3-k)$ -face f'_S in $\text{wVor}(\mathcal{P})$ realized as $\cap_{p \in S} V_p$.

When \mathcal{P} is unweighted, the same definitions yield the standard (unweighted) Voronoi diagram $\text{Vor}(\mathcal{P})$ and its dual Delaunay triangulation $\text{Del}(\mathcal{P})$.

2.3 Unions of balls

Let \mathcal{B} denote the set of balls corresponding to a set of weighted points \mathcal{P} and define the *union of balls* \mathcal{U} as $\cup \mathcal{B}$. It is quite useful to capture the structure of \mathcal{U} using a combinatorial representation like a simplicial complex [34, 35]. Let f_i denote $V_i \cap \partial B_i$ and \mathcal{F} the collection of all such f_i . Observing that $V_i \cap B_j \subseteq V_i \cap B_i \forall B_i, B_j \in \mathcal{B}$, f_i is equivalently defined as the spherical part of $\partial(V_i \cap B_i)$. Consider also the decomposition of \mathcal{U} by the cells of $\text{wVor}(\mathcal{P})$ into $\mathcal{C}(\mathcal{B}) = \{V_i \cap B_i \mid B_i \in \mathcal{B}\}$. The *weighted α -complex* \mathcal{W} is defined as the *geometric realization* of $\mathcal{N}(\mathcal{C}(\mathcal{B}))$ [35], i.e., $\sigma_T \in \mathcal{W}$ if $\{V_i \cap B_i \mid p_i \in T\} \in \mathcal{N}(\mathcal{C}(\mathcal{B}))$. It is not hard to see that \mathcal{W} is a subcomplex of $\text{wDel}(\mathcal{P})$. To see why \mathcal{W} is relevant, consider its *underlying space*; we create a collection containing the convex hull of each simplex in \mathcal{W} and define the *weighted α -shape* \mathcal{J} as the union of this collection. It turns out that the simplices $\sigma_T \in \mathcal{W}$ contained in $\partial \mathcal{J}$ are dual to the faces of $\partial \mathcal{U}$ defined as $\cap_{i \in T} f_i$. In particular, the *corner vertices* of $\partial \mathcal{U}$ correspond to the 2-simplices or facets of $\partial \mathcal{J}$. In fact, every point $q \in \partial \mathcal{U}$ defined by $\cap_{i \in T_q} f_i$, for $T_q \in \mathcal{B}$ and $k+1 = |T_q|$, witnesses the existence of σ_{T_q} in \mathcal{W} ; the k -simplex σ_{T_q} is said to be *exposed* and $\partial \mathcal{J}$ can be defined directly as the collection of all exposed simplices [34]. Moreover, it is well-known that \mathcal{J} is homotopy-equivalent to \mathcal{U} [35].

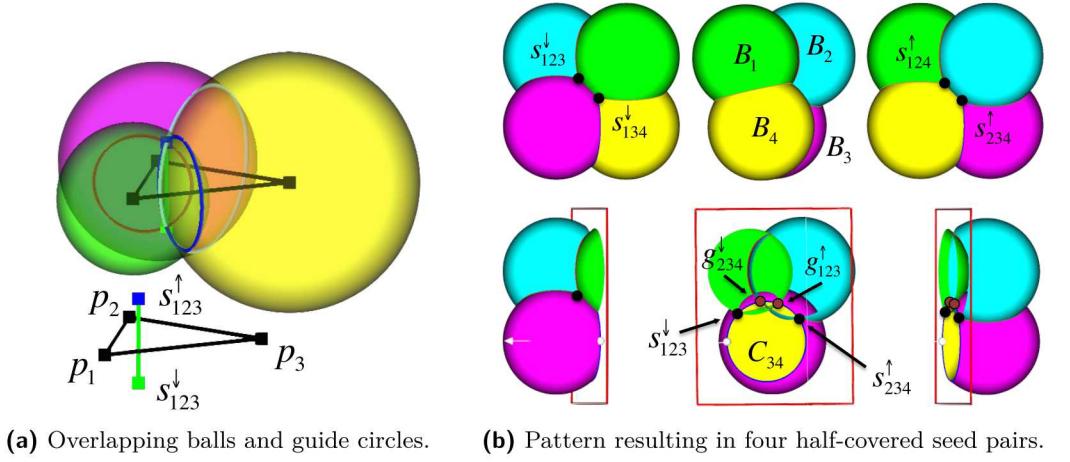
Picking up on that last remark, using unions of balls to approximate an underlying manifold given a set of sample points has been an active subject of study. A union of balls defined using an ϵ -sampling guarantees the approximation of the manifold under suitable conditions on the sampling. Following earlier results on uniform sampling [43], an extension to non-uniform sampling establishes sampling conditions for the isotopic approximation of hypersurfaces and the reconstruction of the medial axis of the manifold [24].

3 Seed placement and surface reconstruction

We determine the location of Voronoi seeds using the union of balls \mathcal{U} . The correctness of our reconstruction depends crucially on how sample balls \mathcal{B} overlap. Assuming a certain structural property on \mathcal{U} , the surface reconstruction is embedded in the dual shape \mathcal{J} .

3.1 Seeds and guides

Central to the method and analysis are triplets of sample spheres, i.e., boundaries of sample balls, corresponding to a *guide triangle* in $\text{wDel}(\mathcal{P})$. The sample spheres associated with the vertices of a guide triangle intersect contributing a pair of *guide points*. The reconstruction consists of Voronoi facets, most of which are guide triangles.



(a) Overlapping balls and guide circles. (b) Pattern resulting in four half-covered seed pairs.

Figure 2 (a) Guide triangle and its dual seed pair. (b) Cutaway view in the plane of circle C_{34} .

When a triplet of spheres $\partial B_i, \partial B_j, \partial B_k$ intersect at exactly two points, the intersection points are denoted by $g_{ijk}^\uparrow = \{g_{ijk}^\uparrow, g_{ijk}^\downarrow\}$ and called a pair of *guide points* or *guides*; see Figure 2a. The associated *guide triangle* t_{ijk} is *dual* to g_{ijk}^\uparrow . We use arrows to distinguish guides on different sides of the manifold with g^\uparrow lying outside \mathcal{O} and g^\downarrow lying inside. We refer to the edges of guide triangles as *guide edges* $e_{ij} = \overline{p_i p_j}$. A guide edge e_{ij} is associated with a dual *guide circle* $C_{ij} = \partial B_i \cap \partial B_j$, as highlighted in Figure 2a.

The Voronoi seeds in $\mathcal{S}^\uparrow \cup \mathcal{S}^\downarrow$ are chosen as the subset of guide points that lie on $\partial \mathcal{U}$. A guide point g which is not interior to any sample ball is *uncovered* and included as a *seed* s into \mathcal{S} ; covered guides are not. We denote *uncovered guides* by s and *covered guides* by g , whenever coverage is known and important. If only one guide point in a pair is covered, then we say the guide pair is *half-covered*. If both guides in a pair are covered, they are ignored. Let $\mathcal{S}_i = \mathcal{S} \cap \partial B_i$ denote the seeds on sample sphere ∂B_i .

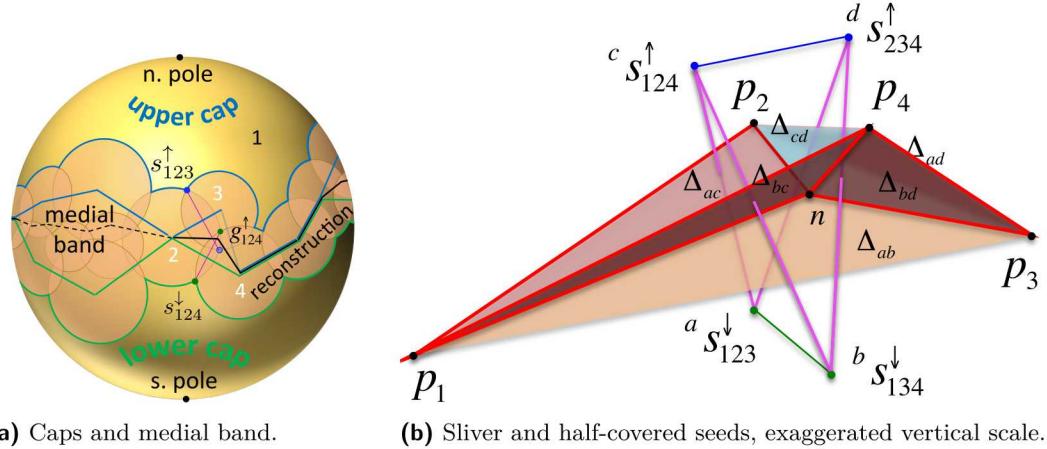
As each guide triangle t_{ijk} is associated with at least one dual seed s_{ijk} , the seed witnesses its inclusion in \mathcal{W} and t_{ijk} is exposed. Hence, t_{ijk} belongs to $\partial \mathcal{J}$ as well. When such t_{ijk} is dual to a single seed s_{ijk} it bounds the interior of \mathcal{J} , i.e., it is a face of a *regular component* of \mathcal{J} ; in the simplest and most common case, t_{ijk} is a facet of a tetrahedron as shown in Figure 3b. When t_{ijk} is dual to a pair of seeds $s_{ijk}^\uparrow, s_{ijk}^\downarrow$, it does not bound the interior of \mathcal{J} and is called a *singular face* of $\partial \mathcal{J}$. All singular faces of $\partial \mathcal{J}$ appear in the reconstructed surface.

3.2 Disk caps

We describe the structural property required on \mathcal{U} along with the consequences exploited by VoroCrust for surface reconstruction. This is partially motivated by the requirement that all sample points on the surface appear as vertices in the output Voronoi mesh.

We define the subset of ∂B_i inside other balls as the *medial band* and say it is *covered*. Let the caps K_i^\uparrow and K_i^\downarrow be the complement of the medial band in the interior and exterior of \mathcal{O} , respectively. Letting n_{p_i} be the normal line through p_i perpendicular to \mathcal{M} , the two intersection points $n_{p_i} \cap \partial B_i$ are called the *poles* of B_i . See Figure 3a.

We require that \mathcal{U} satisfies the following structural property: each ∂B_i has *disk caps*, meaning the medial band is a *topological annulus* and the two caps contain the poles and are *topological disks*. As shown in Figure 3a, all seeds in \mathcal{S}_i^\uparrow and \mathcal{S}_i^\downarrow lie on ∂K_i^\uparrow and ∂K_i^\downarrow , respectively, along the arcs where other sample balls intersect ∂B_i .



(a) Caps and medial band.

(b) Sliver and half-covered seeds, exaggerated vertical scale.

Figure 3 (a) Decomposing the sample sphere ∂B_1 . (b) Uncovered seeds and reconstruction facets. Let $\tau_p \in \text{wDel}(\mathcal{P})$ and $\tau_s \in \text{Del}(\mathcal{S})$ denote the tetrahedra connecting the four samples and the four seeds shown, respectively. s_{123}^\downarrow and s_{134}^\downarrow are the uncovered lower guide seeds, with g_{123}^\uparrow and g_{134}^\uparrow covered. The uncovered upper guide seeds are s_{124}^\uparrow and s_{234}^\uparrow , with g_{124}^\downarrow and g_{234}^\downarrow covered. Δ_{ac} is the Voronoi facet dual to the Delaunay edge between s_{123}^\downarrow and s_{124}^\uparrow , etc. Voronoi facets dual to magenta edges are in the reconstructed surface; those dual to green and blue edges are not. n is the circumcenter of τ_s and appears as a Voronoi vertex in $\text{Vor}(\mathcal{S})$ and a *Steiner vertex* in the surface reconstruction. In general, n is not the circumcenter of the sliver τ_p .

The importance of disk caps is made clear by the following observation. In Section 4, we establish sufficient sampling conditions to ensure \mathcal{U} satisfies this property.

► **Observation 1** (Three upper/lower seeds). *If ∂B_i has disk caps, then each of ∂K_i^\uparrow and ∂K_i^\downarrow has at least three seeds and the seeds on ∂B_i are not all coplanar.*

Proof. Every sphere $S_{j \neq i}$ covers strictly less than one hemisphere of ∂B_i because the poles are uncovered. Hence, each cap is composed of at least three arcs connecting at least three upper seeds $S_i^\uparrow \subset \partial K_i^\uparrow$ and three lower seeds $S_i^\downarrow \subset \partial K_i^\downarrow$. Further, any hemisphere through the poles contains at least one upper and one lower seed. It follows that the set of seeds $S_i = S_i^\uparrow \cup S_i^\downarrow$ is not coplanar. ◀

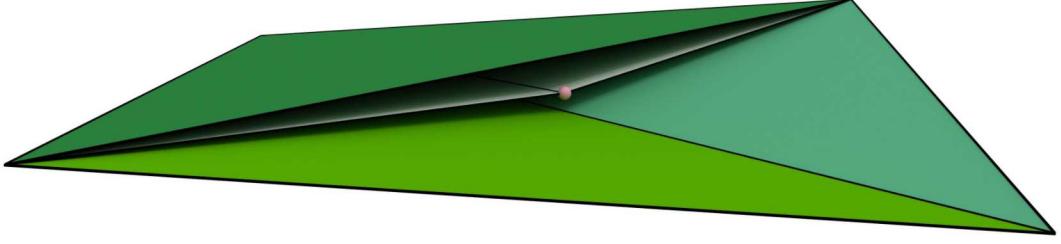
The requirement that all sample points appear as Voronoi vertices follows as a corollary.

► **Corollary 2** (Sample reconstruction). *If ∂B_i has disk caps, then p_i is a vertex in $\text{Vor}(\mathcal{S})$.*

Proof. By Observation 1, the sample is equidistant to at least four seeds which are not all coplanar. It follows that the sample appears as a vertex in the Voronoi diagram and not in the relative interior of a facet or an edge. Being a common vertex to at least one interior and one exterior Voronoi seed, VoroCrust retains this vertex in its output reconstruction. ◀

3.3 Sandwiching the reconstruction in the dual shape of \mathcal{U}

Triangulations of smooth surfaces embedded in \mathbb{R}^3 can have half-covered guides pairs, with one guide covered by the sphere of a fourth sample not in the guide triangle. The tetrahedron formed by the three samples of the guide triangle plus the fourth covering sample is a *sliver*. In this case we do not reconstruct the guide triangle, and also do not reconstruct some guide edges. We show that the reconstructed surface $\hat{\mathcal{M}}$ lies entirely within the region of space bounded by guide triangles, i.e., the α -shape of \mathcal{P} , as stated in the following theorem.



■ **Figure 4** Cutaway view of a sliver tetrahedron $\tau_p \in \mathcal{W} \subseteq \text{wDel}(\mathcal{P})$, drawn to scale. Half-covered guides give rise to the Steiner vertex (pink), which results in a surface reconstruction using four facets (only two are shown) sandwiched within τ . In contrast, filtering $\text{wDel}(\mathcal{P})$ chooses two of the four facets of τ_p , either the bottom two, or the top two (only one is shown).

► **Theorem 3 (Sandwiching).** *If all sample balls have disk caps, then $\hat{\mathcal{M}} \subseteq \mathcal{J}$.*

The simple case of a single isolated sliver tetrahedron is illustrated in Figures 3b, 4 and 2b. A sliver has a pair of lower guide triangles and a pair of upper guide triangles. For instance, t_{124} and t_{234} are the pair of upper triangles in Figure 3b. In such a tetrahedron, there is an edge between each pair of samples corresponding to a non-empty circle of intersection between sample balls, like the circles in Figure 2a. For this circle, the arcs covered by the two other sample balls of the sliver overlap, so each of these balls contributes exactly one uncovered seed, rather than two. In this way the upper guides for the upper triangles are uncovered, but their lower guides are covered; also only the lower guides of the lower triangles are uncovered. The proof of Theorem 3 follows by analyzing the Voronoi cells of the seed points located on the overlapping sample balls and is deferred to Appendix A [1]. Alternatively, Theorem 3 can be seen as a consequence of Theorem 2 in [15].

4 Sampling conditions and approximation guarantees

We take as input a set of points \mathcal{P} sampled from the bounding surface \mathcal{M} such that \mathcal{P} is an ϵ -sample, with $\epsilon \leq 1/500$. We require that \mathcal{P} satisfies the following sparsity condition: for any two points $p_i, p_j \in \mathcal{P}$, $\text{lfs}(p_i) \geq \text{lfs}(p_j) \implies \mathbf{d}(p_i, p_j) \geq \sigma \text{lfs}(p_j)$, with $\sigma \geq 3/4$.

Such a sampling \mathcal{P} can be obtained by known algorithms. Given a suitable representation of \mathcal{M} , the algorithm in [19] computes a loose ϵ' -sample E which is a $\epsilon'(1+8.5\epsilon')$ -sample. More specifically, whenever the algorithm inserts a new sample p into the set E , $\mathbf{d}(p, E) \geq \epsilon' \text{lfs}(p)$. To obtain E as an ϵ -sample, we set $\epsilon'(\epsilon) = (\sqrt{34\epsilon + 1} - 1)/17$. Observing that $3\epsilon/4 \leq \epsilon'(\epsilon)$ for $\epsilon \leq 1/500$, the returned ϵ -sample satisfies our required sparsity condition with $\sigma \geq 3/4$.

We start by adapting Theorem 6.2 and Lemma 6.4 from [24] to the setting just described. For $x \in \mathbb{R}^3 \setminus \mathcal{M}$, let $\Gamma(x) = \mathbf{d}(x, \tilde{x})/\text{lfs}(\tilde{x})$, where \tilde{x} is the closest point to x on \mathcal{M} .

► **Corollary 4.** *For an ϵ -sample \mathcal{P} , with $\epsilon \leq 1/20$, the union of balls \mathcal{U} with $\delta = 2\epsilon$ satisfies:*

1. \mathcal{M} is a deformation retract of \mathcal{U} ,
2. $\partial\mathcal{U}$ contains two connected components, each isotopic to \mathcal{M} ,
3. $\Gamma^{-1}([0, a']) \subset U \subset \Gamma^{-1}([0, b'])$, where $a' = \epsilon - 2\epsilon^2$ and $b' \leq 2.5\epsilon$.

Proof. Theorem 6.2 from [24] is stated for balls with radii within $[a, b]$ times the lfs . We set $a = b = \delta$ and use $\epsilon \leq 1/20$ to simplify fractions. This yields the above expressions for $a' = (1 - \epsilon)\delta - \epsilon$ and $b' = \delta/(1 - 2\delta)$. The general condition requires $(1 - a')^2 + (b' - a' + \delta(1 + 2b' - a'))/(1 - \delta) < 1$, as we assume no noise. Plugging in the values of a' and b' , we verify that the inequality holds for the chosen range of ϵ . ◀

Furthermore, we require that each ball $B_i \in \mathcal{B}$ contributes one facet to each side of $\partial\mathcal{U}$. Our sampling conditions ensure that both poles are outside any ball $B_j \in \mathcal{B}$.

► **Lemma 5 (Disk caps).** All balls in \mathcal{B} have disk caps for $\epsilon \leq 0.066$, $\delta = 2\epsilon$ and $\sigma \geq 3/2$.

Proof. Fix a sample p_i and let x be one of the poles of B_i and $B_x = \mathbb{B}(c, \text{lfs}(p_i))$ the tangent ball at p_i with $x \in B_x$. Letting p_j be the closest sample to x in $P \setminus \{p_i\}$, we assume the worst case where $\text{lfs}(p_j) \geq \text{lfs}(p_i)$ and p_j lies on ∂B_x . To simplify the calculations, take $\text{lfs}(p_i) = 1$ and let ℓ denote $\mathbf{d}(p_i, p_j)$. As lfs is 1-Lipschitz, we get $\text{lfs}(p_j) \leq 1 + \ell$. By the law of cosines, $\mathbf{d}(p_j, x)^2 = \mathbf{d}(p_i, p_j)^2 + \mathbf{d}(p_i, x)^2 - 2\mathbf{d}(p_i, p_j)\mathbf{d}(p_i, x)\cos(\phi)$, where $\phi = \angle_{p_j}p_i c$. Letting $\theta = \angle_{p_i}c p_j$, observe that $\cos(\phi) = \sin(\theta/2) = \ell/2$. To enforce $x \notin B_j$, we require $\mathbf{d}(p_j, x) > \delta \text{lfs}(p_j)$, which is equivalent to $\ell^2 + \delta^2 - \delta\ell^2 > \delta^2(1 + \ell)^2$. Simplifying, we get $\ell > 2\delta^2/(1 - \delta - \delta^2)$ where sparsity guarantees $\ell > \sigma\epsilon$. Setting $\sigma\epsilon > 2\delta^2/(1 - \delta - \delta^2)$ we obtain $4\sigma\epsilon^2 + (8 + 2\sigma)\epsilon - \sigma < 0$, which requires $\epsilon < 0.066$ when $\sigma \geq 3/4$. ◀

Corollary 4 together with Lemma 5 imply that each ∂B_i is decomposed into a covered region $\partial B_i \cap \bigcup_{j \neq i} B_j$, the *medial band*, and two uncovered caps $\partial B_i \setminus \bigcup_{j \neq i} B_j$, each containing one pole. Recalling that seeds arise as pairs of intersection points between the boundaries of such balls, we show that seeds can be classified correctly as either inside or outside \mathcal{M} .

► **Corollary 6.** If a seed pair lies on the same side of \mathcal{M} , then at least one seed is covered.

Proof. Fix such a seed pair $\partial B_i \cap \partial B_j \cap \partial B_k$ and recall that $\mathcal{M} \cap \partial B_i$ is contained in the medial band on ∂B_i . Now, assume for contradiction that both seeds are uncovered and lie on the same side of \mathcal{M} . It follows that $B_j \cap B_k$ intersects B_i away from its medial band, a contradiction to Corollary 4. ◀

Corollary 4 guarantees that the medial band of B_i is a superset of $\Gamma^{-1}([0, a']) \cap \partial B_i$, which means that all seeds s_{ijk} are at least $a' \text{lfs}(\tilde{s}_{ijk})$ away from \mathcal{M} . It will be useful to bound the elevation of such seeds above T_{p_i} , the *tangent plane* to \mathcal{M} at p_i .

► **Lemma 7.** For a seed $s \in \partial B_i$, $\theta_s = \angle_{sp_i}s' \geq 29.34^\circ$ and $\theta_s > \frac{1}{2} - 5\epsilon$, where s' is the projection of s on T_{p_i} , implying $\mathbf{d}(s, s') \geq h_s^\perp \delta \text{lfs}(p_i)$, with $h_s^\perp > 0.46$ and $h_s^\perp > \frac{1}{2} - 5\epsilon$.

Proof. Let $\text{lfs}(p_i) = 1$ and $B_s = \mathbb{B}(c, 1)$ be the tangent ball at p_i with $s \notin B_s$; see Figure 5a. Observe that $\mathbf{d}(s, \mathcal{M}) \leq \mathbf{d}(s, x)$, where $x = \overline{sc} \cap \partial B_s$. By the law of cosines, $\mathbf{d}(s, c)^2 = \mathbf{d}(p_i, c)^2 + \mathbf{d}(p_i, s)^2 - 2\mathbf{d}(p_i, c)\mathbf{d}(p_i, s)\cos(\pi/2 + \theta_s) = 1 + \delta^2 + 2\delta\sin(\theta_s)$. We may write¹ $\mathbf{d}(s, c) \leq 1 + \delta^2/2 + \delta\sin(\theta_s)$. It follows that $\mathbf{d}(s, x) \leq \delta^2/2 + \delta\sin(\theta_s)$. As lfs is 1-Lipschitz and $\mathbf{d}(p_i, x) \leq \delta$, we get $1 - \delta \leq \text{lfs}(x) \leq 1 + \delta$. There must exist a sample p_j such that $\mathbf{d}(x, p_j) \leq \epsilon \text{lfs}(x) \leq \epsilon(1 + \delta)$. Similarly, $\text{lfs}(p_j) \geq (1 - \epsilon(1 + \delta))(1 - \delta)$. By the triangle inequality, $\mathbf{d}(s, p_j) \leq \mathbf{d}(s, x) + \mathbf{d}(x, p_j) \leq \delta^2/2 + \delta\sin(\theta_s) + \epsilon(1 + \delta)$. Setting $\mathbf{d}(s, p_j) < \delta(1 - \delta)(1 - \epsilon(1 + \delta))$ implies $\mathbf{d}(s, p_j) < \delta \text{lfs}(p_j)$, which shows that for small values of θ_s , s cannot be a seed and $p_j \neq p_i$. Substituting $\delta = 2\epsilon$, we get $\theta_s \geq \sin^{-1}(2\epsilon^3 - 5\epsilon + 1/2) \geq 29.34^\circ$ and $\theta_s > 1/2 - 5\epsilon$. ◀

We make frequent use of the following bound on the distance between samples.

► **Claim 8.** $B_i \cap B_j \neq \emptyset \implies \mathbf{d}(p_i, p_j) \leq \kappa\delta \cdot \text{lfs}(p_i)$, with $\kappa = 2/(1 - \delta)$ and $\mathbf{d}(p_i, p_j) \geq \kappa_\epsilon \cdot \text{lfs}(p_i)$ with $\kappa_\epsilon = \sigma\epsilon/(1 + \sigma\epsilon)$.

¹ Define $f(u, v) = \sqrt{1 + u^2 + 2uv} - (1 + u^2/2 + uv)$ and observe that $f(u, -u/2) = 0$ is the only critical value of $f(u, .)$. As $\partial^2 f / \partial v^2 \leq 0$ for $(u, v) \in \mathbb{R} \times [-1, 1]$, we get that $f(u, v) \leq 0$ in this range.

Proof. The upper bound comes from $\mathbf{d}(p_i, p_j) \leq r_i + r_j$ and $\text{lfs}(p_j) \leq \text{lfs}(p_i) + \mathbf{d}(p_i, d_j)$ by 1-Lipschitz, and the lower bound from $\text{lfs}(p_i) - \mathbf{d}(p_i, d_j) \leq \text{lfs}(p_j)$ and the sparsity. \blacktriangleleft

Bounding the circumradii is the culprit behind why we need such small values of ϵ .

► **Lemma 9.** The circumradius of a guide triangle t_{ijk} is at most $\varrho_f \cdot \delta \text{lfs}(p_i)$, where $\varrho_f < 1.38$, and at most $\bar{\varrho}_f \mathbf{d}(p_i, p_j)$ where $\bar{\varrho}_f < 3.68$.

Proof. Let p_i and p_j be the triangle vertices with the smallest and largest lfs values, respectively. From Claim 8, we get $\mathbf{d}(p_i, p_j) \leq \kappa \delta \text{lfs}(p_i)$. It follows that $\text{lfs}(p_j) \leq (1 + \kappa \delta) \text{lfs}(p_i)$. As t_{ijk} is a guide triangle, we know that it has a pair of intersection points $\partial B_i \cap \partial B_j \cap \partial B_k$. Clearly, the seed is no farther than $\delta \text{lfs}(p_j)$ from any vertex of t_{ijk} and the orthoradius of t_{ijk} cannot be bigger than this distance.

Recall that the weight w_i associated with p_i is $\delta^2 \text{lfs}(p_i)^2$. We shift the weights of all the vertices of t_{ijk} by the lowest weight w_i , which does not change the orthocenter. With that $w_j - w_i = \delta^2 (\text{lfs}(p_j)^2 - \text{lfs}(p_i)^2) \leq \delta^2 \text{lfs}(p_i)^2 ((1 + \kappa \delta)^2 - 1) = \kappa \delta^3 \text{lfs}(p_i)^2 (\kappa \delta + 2)$. On the other hand, sparsity ensures that the closest vertex in t_{ijk} to p_j is at distance at least $N(p_j) \geq \sigma \text{lfs}(p_j) \geq \sigma \epsilon (1 - \kappa \delta) \text{lfs}(p_i)$. Ensuring $\alpha^2 \leq (w_j - w_i) / N(p_j)^2 \leq \kappa \delta^3 (2 + \kappa \delta) / (\sigma^2 \epsilon^2 (1 - \kappa \delta)^2) \leq 1/4$ suffices to bound the circumradius of t_{ijk} by $c_{rad} = 1/\sqrt{1 - 4\alpha^2}$ times its orthoradius, as required by Claim 4 in [25]. Substituting $\delta = 2\epsilon$ and $\sigma \geq 3/4$ we get $\alpha^2 \leq 78.97\epsilon$, which corresponds to $c_{rad} < 1.37$. It follows that the circumradius is at most $c_{rad} \delta \text{lfs}(p_j) \leq c_{rad} (1 + \kappa \delta) \delta \text{lfs}(p_i) < 1.38 \delta \text{lfs}(p_i)$.

For the second statement, observe that $\text{lfs}(p_i) \geq (1 - \kappa \delta) \text{lfs}(p_j)$ and the sparsity condition ensures that the shortest edge length is at least $\sigma \text{lfs}(p_i) \geq \sigma \epsilon (1 - \kappa \delta) \text{lfs}(p_j)$. It follows that the circumradius is at most $\frac{\delta c_{rad}}{\sigma \epsilon (1 - \kappa \delta)} < 3.68$ times the length of any edge of t_{ijk} . \blacktriangleleft

Given the bound on the circumradii, we are able to bound the deviation of normals.

► **Lemma 10.** If t_{ijk} is a guide triangle, then (1) $\angle_a(n_{p_i}, n_{p_j}) \leq \eta_s \delta < 0.47^\circ$, with $\eta_s < 2.03$, and (2) $\angle_a(n_t, n_{p_i}) \leq \eta_t \delta < 1.52^\circ$, with $\eta_t < 6.6$, where n_{p_i} is the line normal to \mathcal{M} at p_i and n_t is the normal to t_{ijk} . In particular, t_{ijk} makes an angle at most $\eta_t \delta$ with T_{p_i} .

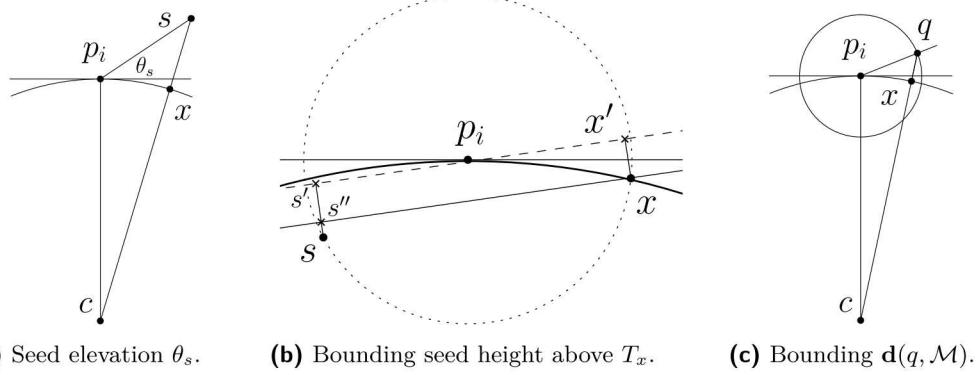
Proof. Claim 8 implies $\mathbf{d}(p_i, p_j) \leq \kappa \delta \text{lfs}(p_i)$ and (1) follows from the Normal Variation Lemma [14] with $\rho = \kappa \delta < 1/3$ yielding $\angle_a(n_{p_i}, n_{p_j}) \leq \kappa \delta / (1 - \kappa \delta)$. Letting R_t denote the circumradius of t , Lemma 9 implies that the $R_t \leq \varrho_f \cdot \delta \text{lfs}(p_i) \leq \text{lfs}(p_i) / \sqrt{2}$ and the Triangle Normal Lemma [29] implies $\angle_a(n_{p^*}, n_t) < 4.57\delta < 1.05^\circ$, where p^* is the vertex of t subtending a maximal angle in t . Hence, $\angle_a(n_{p_i}, n_t) \leq \angle_a(n_{p_i}, n_{p^*}) + \angle_a(n_{p^*}, n_t)$. \blacktriangleleft

Towards establishing homeomorphism, the next lemma on the monotonicity of distance to the nearest seed is critical. First, we show that the nearest seeds to any surface point $x \in \mathcal{M}$ are generated by nearby samples.

► **Lemma 11.** The nearest seed to $x \in \mathcal{M}$ lies on some ∂B_i where $\mathbf{d}(x, p_i) \leq 5.03 \epsilon \text{lfs}(x)$. Consequently, $\mathbf{d}(x, p_i) \leq 5.08 \epsilon \text{lfs}(x)$.

Proof. In an ϵ -sampling, there exists a p_a such that $\mathbf{d}(x, p_a) \leq \epsilon \text{lfs}(x)$, where $\text{lfs}(p_a) \leq (1 + \epsilon) \text{lfs}(x)$. The sampling conditions also guarantee that there exists at least one seed s_a on ∂B_a . By the triangle inequality, we get that $\mathbf{d}(x, s_a) \leq \mathbf{d}(x, p_a) + \mathbf{d}(p_a, s_a) \leq \epsilon \text{lfs}(x) + \delta \text{lfs}(p_a) \leq \epsilon (1 + 2(1 + \epsilon)) \text{lfs}(x) = \epsilon (2\epsilon + 3) \text{lfs}(x)$.

We aim to bound ℓ to ensure $\forall p_i$ s.t. $\mathbf{d}(x, p_i) = \ell \cdot \epsilon \text{lfs}(x)$, the nearest seed to x cannot lie on B_i . Note that in this case, $(1 - \ell\epsilon) \text{lfs}(x) \leq \text{lfs}(p_i) \leq (1 + \ell\epsilon) \text{lfs}(x)$. Let s_i be any seed on B_i . It follows that $\mathbf{d}(x, s_i) \geq \mathbf{d}(x, p_i) - \mathbf{d}(p_i, s_i) \geq \ell \cdot \epsilon \text{lfs}(x) - 2\epsilon \text{lfs}(p_i) \geq \epsilon ((1 - 2\epsilon)\ell - 2) \text{lfs}(x)$.



■ **Figure 5** Constructions used for (a) Lemma 7, (b) Lemma 12 and (c) Theorem 13.

Setting $\epsilon((1-2\epsilon)\ell-2)\text{lfs}(x) \geq \epsilon(2\epsilon+3)\text{lfs}(x)$ suffices to ensure $\mathbf{d}(x, s_i) \geq \mathbf{d}(x, s_a)$, and we get $\ell \geq (2\epsilon+5)/(1-2\epsilon)$. Conversely, if the nearest seed to x lies on B_i , it must be the case that $\mathbf{d}(x, p_i) \leq \ell\text{lfs}(x)$. We verify that $\ell\epsilon = \epsilon(2\epsilon+5)/(1-2\epsilon) < 1$ for any $\epsilon < 0.13$. It follows that $\mathbf{d}(x, p_j) \leq \ell\epsilon/(1-\ell\epsilon)\text{lfs}(p_i)$. ◀

► **Lemma 12.** For any normal segment N_x issued from $x \in \mathcal{M}$, the distance to \mathcal{S}^\uparrow is either strictly increasing or strictly decreasing along $\Gamma^{-1}([0, 0.96\epsilon]) \cap N_x$. The same holds for \mathcal{S}^\downarrow .

Proof. Let n_x be the outward normal and T_x be the tangent plane to \mathcal{M} at x . By Lemma 11, the nearest seeds to x are generated by nearby samples. Fix one such nearby sample p_i . For all possible locations of a seed $s \in \mathcal{S}^\uparrow \cap \partial B_i$, we will show a sufficiently large lower bound on $\langle s - s'', n_x \rangle$, where s'' the projection of s onto T_x .

Take $\text{lfs}(p_i) = 1$ and let $B_s = \mathbb{B}(c, 1)$ be the tangent ball to \mathcal{M} at p_i with $s \in B_s$. Let A be the plane containing $\{p_i, s, x\}$. Assume in the worst case that $A \perp T_{p_i}$ and x is as far as possible from p_i on $\partial B_s \cap T_{p_i}$. By Lemma 11, $\mathbf{d}(p_i, x) \leq 5.08\epsilon$ and it follows that $\theta_x = \angle(n_x, n_{p_i}) \leq 5.08\epsilon/(1-5.08\epsilon) \leq 5.14\epsilon$. This means that T_x is confined within a $(\pi/2 - \theta_x)$ -cocone centered at x . Assume in the worst case that n_x is parallel to A and T_x is tilted to minimize $\mathbf{d}(s, s'')$; see Figure 5b.

Let T'_x be a translation of T_x such that $p_i \in T'_x$ and denote by x' and s' the projections of x and s , respectively, onto T'_x . Observe that T'_x makes an angle θ_x with T_{p_i} . From the isosceles triangle $\triangle p_i c x$, we get that $\theta'_x \leq 1/2 \angle p_i c x = \sin^{-1} 5.08\epsilon/2 \leq 2.54\epsilon$. Now, consider $\triangle p_i x x'$ and let $\phi = \angle x p_i x'$. We have that $\phi = \theta_x + \theta'_x \leq 2.54\epsilon + \delta/(1-\delta) \leq 4.55\epsilon$. Hence, $\sin(\phi) \leq 4.55\epsilon$ and $\mathbf{d}(x, x') \leq 5.08\epsilon \sin(\phi) \leq 0.05\epsilon$. On the other hand, we have that $\angle s p_i s' = \psi \geq \theta_s - \theta_x$ and $\mathbf{d}(s, s') \geq \delta \sin \psi$, where $\theta_s \geq 1/2 - 5\epsilon$ by Lemma 7. Simplifying we get $\sin(\psi) \geq 1/2 - 10.08\epsilon$. The proof follows by evaluating $\mathbf{d}(s, s'') = \mathbf{d}(s, s') - \mathbf{d}(x, x')$. ◀

► **Theorem 13.** For every $x \in \mathcal{M}$ with closest point $q \in \hat{\mathcal{M}}$, and for every $q \in \hat{\mathcal{M}}$ with closest point $x \in \mathcal{M}$, we have $\|xq\| < h_t \epsilon^2 \text{lfs}(x)$, where $h_t < 30.52$. For $\epsilon < 1/500$, $h_t \epsilon^2 < 0.0002$. Moreover, the restriction of the mapping π to $\hat{\mathcal{M}}$ is a homeomorphism and $\hat{\mathcal{M}}$ and \mathcal{M} are ambient isotopic. Consequently, $\hat{\mathcal{O}}$ is ambient isotopic to \mathcal{O} as well.

Proof. Fix a sample $p_i \in \mathcal{P}$ and a surface point $x \in \mathcal{M} \cap B_i$. We consider two cocones centered at x : a p -cocone contains all nearby surface points and a q -cocone contains all guide triangles incident at p_i . By Theorem 3, all reconstruction facets generated by seeds on B_i are sandwiched in the q -cocone.

Lemma 10 readily provides a bound on the q -cocone angle as $\gamma \leq \eta_t \delta$. In addition, since $\mathbf{d}(p_i, x) \leq \delta \text{lfs}(p_i)$, we can bound the p -cocone angle as $\theta \leq 2 \sin^{-1}(\delta/2)$ by Lemma 2 in [7]. We utilize a mixed pq -cocone with angle $\omega = \gamma/2 + \theta/2$, obtained by gluing the lower half of the p -cocone with the upper half of the q -cocone.

Let $q \in \hat{\mathcal{M}}$ and consider its closest point $x \in \mathcal{M}$. Again, fix $p_i \in \mathcal{P}$ such that $x \in B_i$; see Figure 5c. By sandwiching, we know that any ray through q intersects at least one guide triangle, in some point y , after passing through x . Let us assume the worst case that y lies on the upper boundary of the pq -cocone. Then, $\mathbf{d}(q, x) \leq \mathbf{d}(y, y') = h = \delta \sin(\omega) \text{lfs}(p_i)$, where y' is the closest point on the lower boundary of the pq -cocone point to q . We also have that, $\mathbf{d}(p_i, x) \leq \cos(\omega) \delta \text{lfs}(p_i) \leq \delta \text{lfs}(p_i)$, and since lfs is 1-Lipschitz, $\text{lfs}(p_i) \leq \text{lfs}(x)/(1 - \delta)$. Simplifying, we write $\mathbf{d}(q, x) < \delta \omega / (1 - \delta) \cdot \text{lfs}(x) < h_t \epsilon^2 \text{lfs}(x)$.

With $\mathbf{d}(q, x) \leq 0.55 \epsilon \text{lfs}(x)$, Lemma 12 shows that the normal line from any $p \in \mathcal{M}$ intersects $\hat{\mathcal{M}}$ exactly once close to the surface. It follows that for every point $x \in \mathcal{M}$ with closest point $q \in \hat{\mathcal{M}}$, we have $\mathbf{d}(x, q) \leq \mathbf{d}(x, q')$ where $q' \in \hat{\mathcal{M}}$ with x its closest point in \mathcal{M} . Hence, $\mathbf{d}(x, q) \leq h_t \epsilon^2 \text{lfs}(x)$ as well.

Building upon Lemma 12, as a point moves along the normal line at x , it is either the case that the distance to \mathcal{S}^\uparrow is decreasing while the distance to \mathcal{S}^\downarrow is increasing or the other way around. It follows that these two distances become equal at exactly one point on the Voronoi facet above or below x separating some seed $s^\uparrow \in \mathcal{S}^\uparrow$ from another seed $s^\downarrow \in \mathcal{S}^\downarrow$. Hence, the restriction of the mapping π to $\hat{\mathcal{M}}$ is a homeomorphism.

This shows that $\hat{\mathcal{M}}$ and \mathcal{M} homeomorphic. Recall that Corollary 4(3) implies \mathcal{U} is a *topological thickening* [23] of \mathcal{M} . In addition, Theorem 3 guarantees that $\hat{\mathcal{M}}$ is embedded in the interior of \mathcal{U} , such that it separates the two surfaces comprising $\partial \mathcal{U}$. These three properties imply $\hat{\mathcal{M}}$ is isotopic to \mathcal{M} in \mathcal{U} by virtue of Theorem 2.1 in [23]. Finally, as $\hat{\mathcal{M}}$ is the boundary of $\hat{\mathcal{O}}$ by definition, it follows that $\hat{\mathcal{O}}$ is isotopic to \mathcal{O} as well. \blacktriangleleft

5 Quality guarantees and output size

Building upon the analysis in Section 4, we establish a number of quality guarantees on the output mesh. The main result is an upper bound on the *fatness* of all Voronoi cell, i.e., the outradius to inradius ratio. The outradius is the radius of the smallest enclosing ball, and the inradius is the radius of the largest enclosed ball. See Appendix B for the proofs [1].

► **Corollary 14** (Seed height). *If t_{ijk} is a guide triangle with associated seed s , then $\angle s p_i s'' \geq \frac{1}{2} - \eta'_t \epsilon$, where s'' is the projection of s on the plane of t_{ijk} and $\eta'_t \leq 5 + 2\eta_t < 18.18$, implying $\mathbf{d}(s, s'') \geq \hat{h}_s \delta \text{lfs}(p_i)$ with $\hat{h}_s \geq \frac{1}{2} - \eta'_t \epsilon$.*

► **Lemma 15.** For a guide triangle t_{ijk} : (1) edge length ratios are bounded: $\ell_k / \ell_j \leq \kappa_\ell = \frac{2\delta}{1-\delta} \frac{\sigma\epsilon}{1+\sigma\epsilon}$. (2) angles are bounded: $\sin(\theta_i) \geq 1/(2\bar{\varrho}_f)$ implying $\theta_i \in (7.8^\circ, 165^\circ)$. (3) altitudes are bounded: the altitude above e is at least $\alpha_t |e|$, where $\alpha_t = 1/4\bar{\varrho}_f > 0.067$.

Observe that a guide triangle is contained in the Voronoi cell of its seed, even when one of the guides is covered. Hence, the tetrahedron formed by the triangle together with its seed lies inside the cell, and the cell inradius is at least the tetrahedron inradius. Combining the good triangle quality from Lemma 15 with the minimum seed height from Corollary 14, we are able to show a lower bound on the tetrahedron inradius.

To get an upper bound on 3D cell outradius, we must first generate seeds interior to \mathcal{O} . We extend lfs beyond \mathcal{M} , using the point-wise maximal 1-Lipschitz extension [41]: $\text{lfs}(x) = \inf_{p \in \mathcal{M}} (\text{lfs}(p) + \mathbf{d}(x, p))$.

We consider a simple algorithm based on a standard octree over \mathcal{O} . A box is refined if $r > \delta \text{lfs}(c)$, where r is the box radius (half the diagonal) and c is the box center. After refinement terminates, we add an interior seed at c of each empty box, and do nothing with boxes that already contain one or more guide seeds. Box sizes are naturally balanced and slowly varying, and the lfs at any point in the box is within constant factors of $\text{lfs}(c)$. Applying this scheme, we obtain the following.

- **Lemma 16.** The aspect ratio of interior cells is at most $\frac{8\sqrt{3}(1+\delta)}{1-3\delta} < 14.1$.
- **Lemma 17.** The aspect ratio of boundary cells is at most $\frac{4(1+\delta)}{(1-3\delta)(1-\delta)^2 \varrho_v} < 13.65$.

Armed with the aspect ratio bounds, we proceed to bound the output size, the number of seeds (cells), in terms of lfs. The integral of $1/\text{lfs}^3$ over a single cell is bounded above by a constant, because the cell inradius and outradius are both bounded by lfs, and lfs is 1-Lipschitz. Thus, the integral of $1/\text{lfs}^3$ over \mathcal{O} in effect counts the cells.

- **Lemma 18.** $|\mathcal{S}| \leq 18\sqrt{3}/\pi \cdot \epsilon^{-3} \int_{\mathcal{O}} \text{lfs}^{-3}$.

6 Conclusions

We have analyzed an abstract version of the VoroCrust algorithm for volumes bounded by smooth surfaces. We established several guarantees on its output, provided the input samples satisfy certain conditions. In particular, the reconstruction is isotopic to the underlying surface and all 3D Voronoi cells have bounded fatness, i.e., outradius to inradius ratio. The triangular faces of the reconstruction have bounded angles and edge-length ratios, except perhaps in the presence of slivers. In a forthcoming paper [3], we describe the design and implementation of the complete VoroCrust algorithm, which generates conforming Voronoi meshes of realistic models, possibly containing sharp features, and produces samples that follow a natural sizing function and ensure output quality.

For future work, it would be interesting to ensure both guides are uncovered, or both covered. This might be achievable by additional conditions on the ϵ -sampling, or a different finite sampling algorithm. The significance would be that no tetrahedral slivers arise and no Steiner points are introduced. Further, the surface reconstruction would be composed entirely of guide triangles, so it would be easy to show that triangle normals converge to surface normals as sample density increases. Alternatively, where Steiner points are introduced on the surface, it would be helpful to have conditions that guaranteed the triangles containing Steiner points have good quality. In addition, the minimum edge length in a Voronoi cell can be a limiting factor in certain numerical solvers. Post-processing by mesh optimization techniques [5] can help eliminate short Voronoi edges away from the surface. Finally, we expect that the abstract algorithm analyzed in this paper can be extended to higher dimensions.

Acknowledgements. We thank Tamal Dey for helpful discussions regarding the homeomorphic reconstruction literature. The research of Chandrajit Bajaj was supported in part by a contract from Sandia, #1439100, and a grant from NIH, R01-GM117594. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research (ASCR), Applied Mathematics Program. Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

References

- 1 Ahmed Abdelkader, Chandrajit L. Bajaj, Mohamed S. Ebeida, Ahmed H. Mahmoud, Scott A. Mitchell, John D. Owens, and Ahmad A. Rushdi. Sampling Conditions for Conforming Voronoi Meshing by the VoroCrust Algorithm. *CoRR*, arXiv:1803.06078, 2018. URL: <http://arxiv.org/abs/1803.06078>.
- 2 Ahmed Abdelkader, Chandrajit L. Bajaj, Mohamed S. Ebeida, Ahmed H. Mahmoud, Scott A. Mitchell, John D. Owens, and Ahmad A. Rushdi. VoroCrust Illustrated: Theory and Challenges (Multimedia Contribution). In *34th International Symposium on Computational Geometry (SoCG 2018)*, 2018.
- 3 Ahmed Abdelkader, Chandrajit L. Bajaj, Mohamed S. Ebeida, Ahmed H. Mahmoud, Scott A. Mitchell, John D. Owens, and Ahmad A. Rushdi. VoroCrust: Voronoi meshing without clipping. *in preparation*, 2018.
- 4 Ahmed Abdelkader, Chandrajit L. Bajaj, Mohamed S. Ebeida, and Scott A. Mitchell. A seed placement strategy for conforming Voronoi meshing. In *29th Canadian Conference on Computational Geometry, CCCG 2017*, pages 95–100, 2017.
- 5 Ahmed Abdelkader, Ahmed H. Mahmoud, Ahmad A. Rushdi, Scott A. Mitchell, John D. Owens, and Mohamed S. Ebeida. A constrained resampling strategy for mesh improvement. *Computer Graphics Forum*, 36(5):189–201, 2017.
- 6 O. Aichholzer, F. Aurenhammer, B. Kornberger, S. Plantinga, G. Rote, A. Sturm, and G. Vegter. Recovering structure from r-sampled objects. *Computer Graphics Forum*, 28(5):1349–1360, 2009.
- 7 Nina Amenta and Marshall Bern. Surface reconstruction by Voronoi filtering. *Discrete & Computational Geometry*, 22(4):481–504, December 1999.
- 8 Nina Amenta, Marshall Bern, and David Eppstein. The crust and the β -skeleton: Combinatorial curve reconstruction. *Graphical models and image processing*, 60(2):125–135, 1998.
- 9 Nina Amenta, Marshall Bern, and David Eppstein. Optimal point placement for mesh smoothing. *Journal of Algorithms*, 30(2):302–322, 1999.
- 10 Nina Amenta, Marshall Bern, and Manolis Kamvysselis. A new Voronoi-based surface reconstruction algorithm. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '98*, pages 415–421. ACM, 1998.
- 11 Nina Amenta, Sunghee Choi, Tamal K Dey, and Naveen Leekha. A simple algorithm for homeomorphic surface reconstruction. In *16th Annual Symposium on Computational Geometry*, pages 213–222. ACM, 2000.
- 12 Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. The power crust. In *Proceedings of the Sixth ACM Symp. on Solid Modeling and Applications, SMA '01*, pages 249–266, 2001.
- 13 Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. The power crust, unions of balls, and the medial axis transform. *Computational Geometry*, 19(2):127–153, July 2001.
- 14 Nina Amenta and Tamal K. Dey. Normal variation for adaptive feature size. *CoRR*, abs/1408.0314, 2014.
- 15 Nina Amenta and Ravi Krishna Kolluri. The medial axis of a union of balls. *Computational Geometry*, 20(1):25 – 37, 2001. Selected papers from the 12th Annual Canadian Conference.
- 16 Nina Amenta, Thomas J. Peters, and Alexander C. Russell. Computational topology: ambient isotopic approximation of 2-manifolds. *Theoretical Computer Science*, 305(1):3 – 15, 2003. Topology in Computer Science.
- 17 L. Beirão da Veiga, F. Brezzi, L. D. Marini, and A. Russo. The hitchhiker’s guide to the virtual element method. *Mathematical Models and Methods in Applied Sciences*, 24(08):1541–1573, 2014.

- 18 Joseph E. Bishop. Simulating the pervasive fracture of materials and structures using randomly close packed Voronoi tessellations. *Computational Mechanics*, 44(4):455–471, September 2009.
- 19 Jean-Daniel Boissonnat and Steve Oudot. Provably good sampling and meshing of surfaces. *Graphical Models*, 67(5):405 – 451, 2005. Solid Modeling and Applications.
- 20 Tyson Brochu, Christopher Batty, and Robert Bridson. Matching fluid simulation elements to surface geometry and topology. *ACM Trans. Graph.*, 29(4):47:1–47:9, July 2010.
- 21 F. Cazals, T. Dreyfus, S. Sachdeva, and N. Shah. Greedy geometric algorithms for collection of balls, with applications to geometric approximation and molecular coarse-graining. *Computer Graphics Forum*, 33(6):1–17, 2014.
- 22 Frederic Cazals, Harshad Kanhere, and Sébastien Loriot. Computing the volume of a union of balls: A certified algorithm. *ACM Trans. Math. Softw.*, 38(1):3:1–3:20, December 2011.
- 23 Frédéric Chazal and David Cohen-Steiner. A condition for isotopic approximation. *Graphical Models*, 67(5):390 – 404, 2005. Solid Modeling and Applications.
- 24 Frédéric Chazal and André Lieutier. Smooth manifold reconstruction from noisy and non-uniform approximation with guarantees. *Computational Geometry*, 40(2):156 – 170, 2008.
- 25 Siu-Wing Cheng, Tamal K. Dey, Herbert Edelsbrunner, Michael A. Facello, and Shang-Hua Teng. Silver exudation. *J. ACM*, 47(5):883–904, September 2000.
- 26 Siu-Wing Cheng, Tamal K Dey, and Jonathan Shewchuk. *Delaunay Mesh Generation*. CRC Press, 2012.
- 27 David Cohen-Steiner, Éric Colin de Verdière, and Mariette Yvinec. Conforming Delaunay triangulations in 3D. In *Proceedings of the Eighteenth Annual Symposium on Computational Geometry*, SCG '02, pages 199–208, 2002.
- 28 Kyle Sykes David Letscher. On the stability of medial axis of a union of disks in the plane. In *28th Canadian Conference on Computational Geometry, CCCG 2016*, pages 29–33, 2016.
- 29 Tamal K. Dey. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis*. Cambridge University Press, New York, NY, USA, 2006.
- 30 Tamal K Dey, Kuiyu Li, Edgar A Ramos, and Raphael Wenger. Isotopic reconstruction of surfaces with boundaries. In *Computer Graphics Forum*, volume 28:5, pages 1371–1382. Wiley Online Library, 2009.
- 31 Tamal K. Dey and Lei Wang. Voronoi-based feature curves extraction for sampled singular surfaces. *Computers & Graphics*, 37(6):659–668, October 2013. Shape Modeling International (SMI) Conference 2013.
- 32 Liyun Duan and Florent Lafarge. Image partitioning into convex polygons. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3119–3127, June 2015.
- 33 Mohamed S. Ebeida and Scott A. Mitchell. Uniform random Voronoi meshes. In *International Meshing Roundtable (IMR)*, pages 258–275, 2011.
- 34 Herbert Edelsbrunner. *Weighted alpha shapes*. University of Illinois at Urbana-Champaign, Department of Computer Science, 1992.
- 35 Herbert Edelsbrunner. The union of balls and its dual shape. *Discrete & Computational Geometry*, 13(3):415–440, Jun 1995.
- 36 Herbert Edelsbrunner and Ernst Peter Mücke. Simulation of simplicity: A technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Graph.*, 9(1):66–104, January 1990.
- 37 Robert Eymard, Thierry Gallouët, and Raphaële Herbin. Finite volume methods. In *Techniques of Scientific Computing (Part 3)*, volume 7 of *Handbook of Numerical Analysis*, pages 713 – 1018. Elsevier, 2000.

38 Ø S. Klemetsdal, R. L. Berge, K.-A. Lie, H. M. Nilsen, and O. Møyner. *SPE-182666-MS*, chapter Unstructured Gridding and Consistent Discretizations for Reservoirs with Faults and Complex Wells. Society of Petroleum Engineers, 2017.

39 Dmitri Kuzmin. A guide to numerical methods for transport equations. *University Erlangen-Nuremberg*, 2010.

40 Romain Merland, Guillaume Caumon, Bruno Lévy, and Pauline Collon-Drouaillet. Voronoi grids conforming to 3D structural features. *Computational Geosciences*, 18(3):373–383, Aug 2014.

41 Gary L. Miller, Dafna Talmor, and Shang-Hua Teng. Data generation for geometric algorithms on non-uniform distributions. *International Journal of Computational Geometry and Applications*, 09(06):577–597, 1999.

42 Michael Murphy, David M. Mount, and Carl W. Gable. A point-placement strategy for conforming Delaunay tetrahedralization. *International Journal of Computational Geometry & Applications*, 11(06), 2001.

43 Partha Niyogi, Stephen Smale, and Shmuel Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete & Computational Geometry*, 39(1):419–441, Mar 2008.

44 Atsuyuki Okabe, Barry Boots, Kokichi Sugihara, and Sung Nok Chiu. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, volume 501. John Wiley & Sons, 2009.

45 Alexander Rand and Noel Walkington. Collars and intestines: Practical conforming Delaunay refinement. In *Proceedings of the 18th International Meshing Roundtable*, 2009.

46 Chris Rycroft. Voro++: A three-dimensional Voronoi cell library in C++. *Chaos*, 19(4), 2009. Software available online at <http://math.lbl.gov/voro++/>.

47 Manuel Lorenzo Sents and Carl W. Gable. Coupling LaGrit Unstructured Mesh Generation and Model Setup with TOUGH2 Flow and Transport. *Comput. Geosci.*, 108(C):42–49, November 2017.

48 H. Si, K. Gärtner, and J. Fuhrmann. Boundary conforming Delaunay mesh generation. *Computational Mathematics and Mathematical Physics*, 50(1):38–53, Jan 2010.

49 Dong-Ming Yan, Bruno Lévy, Yang Liu, Feng Sun, and Wenping Wang. Isotropic remeshing with fast and exact computation of restricted Voronoi diagram. *Computer Graphics Forum*, 28(5):1445–1454, July 2009.