

Computing Medial Axis Transform with Feature Preservation via Restricted Power Diagram

NINGNA WANG, University of Texas at Dallas, USA BIN WANG, Tsinghua University, China WENPING WANG, Texas A&M University, USA XIAOHU GUO*, University of Texas at Dallas, USA

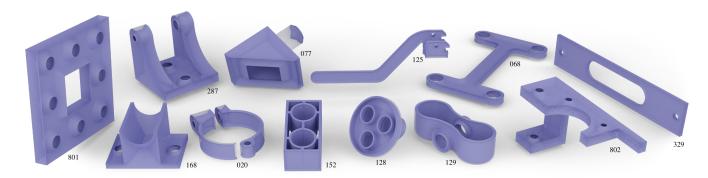


Fig. 1. A gallery of our feature-preserving 3D medial axis results (in violet) from inputs of CAD meshes (in transparency) with sharp features.

We propose a novel framework for computing the medial axis transform of 3D shapes while preserving their medial features via restricted power diagram (RPD). Medial features, including external features such as the sharp edges and corners of the input mesh surface and internal features such as the seams and junctions of medial axis, are important shape descriptors both topologically and geometrically. However, existing medial axis approximation methods fail to capture and preserve them due to the fundamentally undersampling in the vicinity of medial features, and the difficulty to build their correct connections. In this paper we use the RPD of medial spheres and its affiliated structures to help solve these challenges. The dual structure of RPD provides the connectivity of medial spheres. The surfacic restricted power cell (RPC) of each medial sphere provides the tangential surface regions that these spheres have contact with. The connected components (CC) of surfacic RPC give us the classification of each sphere, to be on a medial sheet, a seam, or a junction. They allow us to detect insufficient sphere sampling around medial features and develop necessary conditions to preserve them. Using this RPD-based framework, we are able to construct high quality medial meshes with features preserved. Compared with existing sampling-based or voxel-based methods, our method is the first one that can preserve not only external features but also internal features of medial axes.

*Corresponding author

Authors' addresses: Ningna Wang, University of Texas at Dallas, USA, ningna.wang@utdallas.edu; Bin Wang, Tsinghua University, China, wangbins@tsinghua.edu.cn; Wenping Wang, Texas A&M University, USA, wenping@tamu.edu; Xiaohu Guo, University of Texas at Dallas, USA, xguo@utdallas.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery. 0730-0301/2022/12-ART188 \$15.00 https://doi.org/10.1145/3550454.3555465 CCS Concepts: • Computing methodologies \rightarrow Shape analysis.

Additional Key Words and Phrases: Medial Axis Transform, Feature Preservation, Restricted Power Diagram

ACM Reference Format:

Ningna Wang, Bin Wang, Wenping Wang, and Xiaohu Guo. 2022. Computing Medial Axis Transform with Feature Preservation via Restricted Power Diagram. *ACM Trans. Graph.* 41, 6, Article 188 (December 2022), 18 pages. https://doi.org/10.1145/3550454.3555465

1 INTRODUCTION

The medial axis [Blum et al. 1967] is a fundamental geometric structure and has been widely used in approximating, simplifying, and analyzing shapes. The medial axis \mathcal{M} of a 3D shape \mathcal{S} is simply defined as the set of centers of maximally-inscribed spheres touching two or more points on the surface ∂S . Topologically, the medial axis \mathcal{M} is homotopy equivalent to \mathcal{S} . Geometrically, the medial axis \mathcal{M} captures the protrusions and components of \mathcal{S} . The medial axis transform (MAT) is the combination of the medial axis and the radius function defines on it. We represent MAT using a simplicial complex called medial mesh $\mathcal{M}_{\mathcal{S}}$ and approximate the input 3D shape \mathcal{S} by the union of enveloping volumes of the medial primitives (medial cones and medial slabs, see Sec. 3.1) of the medial mesh.

We have observed that *medial features* play a significant role in guaranteeing many topological and geometric properties of medial axis. *External features* such as sharp edges and corners of the input 3D shape, which are common to CAD models, represent the non-smoothness of the surface (see Fig. 2 (f) lines in black). It follows from the definition that the medial axis passes through the points where the surface is locally convex and non-smooth. *Internal features*, on the other hand, defines the inner topological structure of seams and junctions (see Fig. 2 (f) lines in red) so that medial axis has a

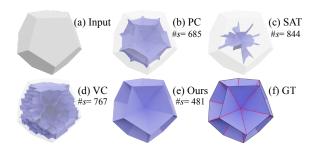


Fig. 2. Approximations of medial axis of (a) a dodecahedron model computed with similar number of medial spheres (#s), using (b) a sampling-based method Power Crust (PC) [Amenta et al. 2001a], (c) another sampling-based method Scale Axis Transform (SAT) [Miklos et al. 2010], (d) a voxel-based method Voxel Core (VC) [Yan et al. 2018], and (e) our method. The ground truth medial axis (f) highlights external features in black and internal features in red.

natural decomposition into multiple manifold sheets. Applications such as hexahedral mesh generation of CAD models [Quadros et al. 2004] [Sampl 2000] relies on internal features as a starting point to perform solid meshing. As a result, the medial features, both external and internal, constitute the foundation of medial axis as skeletal shape descriptors [Tagliasacchi et al. 2016].

However, all existing medial axis approximation methods that are designed to handle large input 3D meshes (e.g., tens of thousands of triangle faces as shown in all models of Fig. 1) fail to preserve external features. For CAD models that commonly contain external features like convex sharp edges and corners, the approximated medial axis usually stops before touching these features, as shown in examples of Fig. 2 (b)(c)(d), in which cases the structures of medial axis are not complete. For all methods depending on inner Voronoi balls of surface samples, no matter how dense the surface sampling points are, the density condition designed based on their local feature size (LFS) [Amenta et al. 2001a] can never be achieved at convex sharp edges and corners. This is because their LFS converges to zero near the external features, causing their sampling density to converge to infinity. Even though weak feature size (WFS) [Chazal and Lieutier 2005a] was proposed to bypass the issue of vanishing LFS for non-smooth shapes, there is no practical method yet to preserve those external features in the resulting medial axis.

Furthermore, internal medial features are even more difficult to capture since they are hidden information of the input shape, as shown in Fig. 2. The essential reason of this failure is because of the under-sampling of medial spheres on internal features (red lines in Fig. 2 (f)). Existing solutions tend to increase surface sampling rate (or decrease voxel size) globally in order to maintain the internal features. This eventually results in a large amount of non-feature medial spheres being sampled without any guarantee to capture the topologically-important internal features.

In order to generate a medial mesh with feature preservation, we need to answer the following three questions: (1) How to sample and update feature spheres to guarantee their tangential properties? For example, spheres on medial seam are tangential to three points on the surface, while spheres on medial junction are tangential to four. (2) How to identify insufficient sampling on medial features,

for both internal and external ones? (3) How to connect all medial spheres with correct topology and geometry? These open questions varies for these two types of medial features:

- External features can be trivially identified on the input mesh surface using the dihedral angle of two incident polygonal faces. We can sample zero-radius medial spheres on those convex external features. However, there are still three major challenges: (a) How to create the connectivity between these zero-radius feature spheres and their nearby non-feature spheres is still an open question. (b) The connectivity of neighboring feature spheres on external sharp edges could be broken by the "invasion" of nearby non-feature medial spheres. (c) Multiple convex and concave sharp edges could meet at a corner, where neighboring edges could form small angles, causing a "corner cap" that connects feature spheres at neighboring sharp edges.
- Internal features, on the other hand, are difficult to identify in the first place. It it known that with more surface samples used to construct inner Voronoi balls, the internal features are more likely to be well-captured. However, there is no indication to tell whether the sphere sampling on internal feature is sufficient or not, especially at the vicinity of external features in CAD models where the sampling is typically not dense enough.

In this paper, we present the first framework for computing an approximated MAT that preserves both external and internal medial features for an input 3D mesh surface while ensuring the approximation accuracy. Our method is based on a novel insight that the surface restricted power diagram (RPD) can provide us auxiliary knowledge about the surface regions which medial spheres have tangential contacts with. This gives us the ability to classify medial spheres based on their contacts to the boundary surface (see details in Sec. 3.3), and sample or update spheres to their ideal position by solving a quadratic energy optimization problem (as answers for question (1) above). We further develop the necessary conditions to preserve both external and internal features based on the information of restricted power cells (RPCs) of these feature spheres with their neighbors (as answers for question (2) above). The restricted regular triangulation (RRT, as dual of RPD) implies the connectivity [Amenta et al. 2001b] of discrete medial spheres to form a mesh approximation of the medial axis (as partial answers for question (3) above). In our experiments as shown in the Supplementary Material, we do encounter some models whose generated medial meshes are not topology preserving. We leave to our future work a theoreticallysound answer to question (3) with topology preservation (see Sec. 7). Fig. 1 shows a gallery of our feature-preserving medial meshes computed from input CAD meshes with sharp features.

The contributions of this paper can be summarized as follows:

- First, we formulate an energy optimization framework (in Sec. 4.2) for all medial spheres of different types, to update their positions and radii as close as possible to ground truth, based on the tangential information derived from their surface RPCs.
- Second, we present a complete RPD-based framework (in Sec. 4.3) for computing 3D medial mesh while preserving

medial features for any input 3D mesh surfaces. We show that the medial mesh derived from RRT (as dual of RPD) could be refined with our geometry-guided thinning algorithm (in Sec. 4.4), in order to keep the thinness property of the medial

• Third, we propose three feature preservation strategies (in Sec. 5) for preserving external edge features, external corner features, and internal features, under the guidance of surface RPCs. Our method not only preserve external features that none of existing methods can work, but also requires fewer medial sphere samples to generate high quality internal features compared with existing methods.

2 RELATED WORKS

In this section, we shall review the representative approaches for MAT computation. For more extensive discussion on medial axes and other medial representations, we refer readers to those survey articles [Siddiqi and Pizer 2008; Tagliasacchi et al. 2016].

Computing 3D exact MAT of polyhedra with feature preservation was explored by earlier algorithms [Culver et al. 2004; Milenkovic 1993; Sherbrooke et al. 1996] using seam tracing. Due to the complexity of these exact computation algorithms, they mainly aim for computing medial axes of a simple class of shapes, e.g., polyhedra composed of up to about 20 faces, which is impractical for more complicated models. Note that all the models shown in this paper have over tens of thousands of faces. Thus a significant branch of later researches resort to compute an approximated MAT (instead of exact computation) in order for them to be suitable for real-world applications [Amenta et al. 2001a,b; Chazal and Lieutier 2005b; Dev and Zhao 2002, 2004; Miklos et al. 2010; Pizer et al. 2003; Saha et al. 2016; Sobiecki et al. 2014].

As medial axis is notorious for its sensitivity to boundary perturbations, a closely related problem is MAT simplification to identify significant and stable parts of the medial axis, and we refer reader to recent excellent works [Dou et al. 2022; Faraj et al. 2013; Li et al. 2015; Pan et al. 2019; Yan et al. 2016]. While all existing MAT simplification methods are designed for smooth shapes, we believe there could be new MAT simplification methods for non-smooth shapes in the future, by following our feature preservation framework of this paper. It is worth noting that many MAT simplification methods require an initial approximation of MAT in order to prune the noisy branches of the medial axis. The method presented in this paper can provide an initial feature-preserved MAT for any future MAT simplification methods for CAD models.

For 2D smooth shape, it has been proved that the subset of Voronoi diagram of boundary samples provides a structure topologically and geometrically converging to the medial axis [Brandt and Algazi 1992]. Unfortunately, this approximation does not hold for 3D smooth shapes due to the existence of "slivers" which are tetrahedra with small volume in Delaunay triangulation of boundary samples. This leads to their Voronoi vertices (centers of circumscribing spheres for sliver tetrahedra) very close to boundary but far away from the medial axis [Amenta et al. 2001b].

Angle-based filtering methods are approaches to filter the Voronoi diagram of obtained boundary samples or other derivative structures to approximate the medial axis. Amenta et al. [2001a] proposed "poles" of Voronoi diagram and show that the power shape converges to the medial axis as the sampling density increases. Several other methods [Brandt and Algazi 1992; Dey and Zhao 2002, 2004] consider a subset of Voronoi diagram of boundary samples that satisfy an angle criteria given a user-specified threshold. These methods are known to be difficult to preserve the topology of the input shape as the filtered subset tends to have many holes and isolated elements.

 λ -medial axis methods [Chazal and Lieutier 2005b; Pizer et al. 2003] use the radius of the closest medial sphere as filtering criteria and discard a medial sphere if its radius is smaller than a given threshold λ . As a result, the medial axis consists of medial spheres such that the smallest enclosing sphere of the nearest boundary sample set has a radius equals or larger than λ .

Voxel-based methods belong to an completely different category that approximate the medial axis of voxel shape by selecting a subset of voxels that share similar properties as medial axis [Saha et al. 2016; Sobiecki et al. 2014]. The state-of-the-art method Voxel Core [Yan et al. 2018] can well approximate the medial axis of any smooth shape while guaranteeing the topological correctness of the generated medial axis, given a voxelization of the shape at sufficiently high resolution. A common drawback of these methods is that they require fine voxel resolutions hence have high computational cost in order to achieve a comparable geometric accuracy as samplingbased methods. Voxel Core requires an additional λ -pruning which shrinks the medial axis while removing ill-posed structure. This pruning operation makes the medial mesh incomplete especially for non-smooth regions like sharp edges or corners.

None of the above MAT approximation methods, however, consider handling 3D shapes with non-smooth regions such as sharp edges and corners (as we call them external features), and none of them consider preserving the internal features of medial axis. To the best of our knowledge, only Dey et al. proposed a remedial method CAD_MEDIAL [Dey et al. 2003] which extends their previous method MEDIAL [Dey and Zhao 2002] to complete the structure of medial axis with external features. But the sampling condition proposed in CAD MEDIAL near sharp edges is too impractical to be achieved, leading to failures to preserve external features in many shapes. Comparing with these methods, our method for approximating the medial axis is the first one that preserves not only external features but also internal features of the medial axis.

PRELIMINARIES

3.1 Medial Axis and Medial Mesh

Given a closed, oriented, and bounded shape $S \in \mathbb{R}^3$, the *medial axis* \mathcal{M} is defined as the locus of centers of spheres that are tangent to two or more points on the boundary of S, or ∂S , without containing any other boundary points in its interior.

The medial axis transform (MAT) is formed by the medial axis M together with its radius function. To approximate the MAT of a 3D shape S, we use a non-manifold medial mesh \mathcal{M}_s consisting of triangles and edges. Each vertex of \mathcal{M}_s represents a medial sphere $\mathbf{m} = (\theta, r)$, where $\theta \in \mathbb{R}^3$ is the sphere center and $r \in \mathbb{R}$ is its radius. MAT can be used to reconstruct the surface through the union of enveloping volumes of its medial primitives [Li et al. 2015]. For

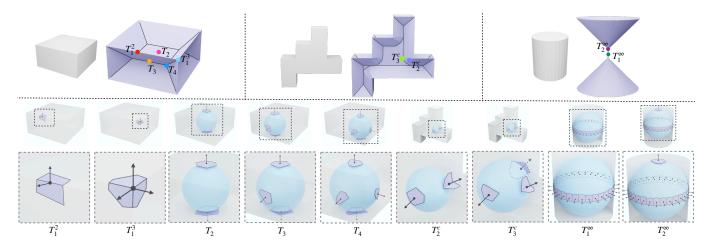


Fig. 3. Examples of different classes of medial spheres in the context of their connected components (CCs) of restricted power cell (RPC). For a medial sphere, the subscript of its class type represents the number of CCs of its RPC, and the superscript represents the number of surface normals for the contact points in each CC. Note that we use the simplified notations here by removing all appearance of "1" from the superscript, *i.e.*, $T_2^{1,1,1}$ is simply written as T_3 .

example, the enveloping volume of an edge of the medial mesh is called a *medial cone*, which is the linear interpolation of two spheres $e_{ij} = t\mathbf{m}_i + (1-t)\mathbf{m}_j, t \in [0,1]$. The enveloping volume of a triangle face f_{ijk} of the medial mesh is called a *medial slab* that is the linear interpolation of three spheres \mathbf{m}_i , \mathbf{m}_j , and \mathbf{m}_k .

In this paper, we assume the boundary surface ∂S is provided as a watertight and manifold 3D triangular mesh, with those sharp feature edges and corners pre-labeled on the mesh. We do not have any other assumption about these sharp features, *e.g.*, how small a dihedral angle is allowed to be on a sharp edge, etc. In fact, the case when multiple sharp edges (convex and/or concave) meet at a corner is particularly challenging, *e.g.*, two adjacent sharp edges forming a small angle around a corner. Our feature preservation strategy is designed to handle all these challenging cases as discussed in Sec. 5.

3.2 Restricted Power Diagram (RPD)

The Voronoi diagram of a set of generators $\{\theta_i\}_{i=1}^n$ is a partition of the domain $\Omega \subset \mathbb{R}^d$ into a set of cells. Each cell Ω_i^{vor} consists of the points $\mathbf{x} \in \Omega$ closest to a particular generator θ_i :

$$\Omega_i^{vor}: \{\mathbf{x} \in \Omega | ||\mathbf{x} - \boldsymbol{\theta}_i|| \le ||\mathbf{x} - \boldsymbol{\theta}_i||, j \ne i\}. \tag{1}$$

The vertices of these cells are called *Voronoi vertices*. The *Voronoi ball* in \mathbb{R}^3 centered at a Voronoi vertex has at least 4 generators on its boundary, and no generator in its interior [Amenta et al. 2001a].

Power diagram [Aurenhammer 1987] is a generalization of Voronoi diagram by weighting the given generators, and coincides with Voronoi diagram in the special case that all points have equal weights. Given a set of weighted generators $\{\mathbf{m}_i = (\theta_i, r_i)\}_{i=1}^n$, a power cell Ω_i^{pow} is defined as:

$$\Omega_i^{pow}: \{\mathbf{x} \in \Omega | d_{pow}(\mathbf{x}, \mathbf{m}_i) \le d_{pow}(\mathbf{x}, \mathbf{m}_j), j \ne i\}, \tag{2}$$

where $d_{pow}(\mathbf{x}, \mathbf{m}_i) = ||\mathbf{x} - \boldsymbol{\theta}_i||^2 - r_i^2$ is the power distance between any point \mathbf{x} and the weighted generator $(\boldsymbol{\theta}_i, r_i)$.

A power diagram restricted within a bounded shape $S \subset \mathbb{R}^n$ is called a *restricted power diagram* (RPD) \mathcal{R} , where all power cells

overlap with the shape S. The RPD consists of a set of *restricted* power cells (RPC) $\{\omega_i\}_{i=1}^n$. Each cell ω_i is the restriction of the power cell Ω_i^{pow} of the weighted generator (θ_i, r_i) within S:

$$\omega_i = \Omega_i^{pow} \cap \mathcal{S}. \tag{3}$$

The surface RPC is the restriction of the power cell within the boundary surface of the shape, denoted as:

$$\partial \omega_i = \Omega_i^{pow} \cap \partial S. \tag{4}$$

Robust and exact computation of surface RPD on a surface represented as triangle mesh is a non-trivial task. Our implementation extends Yan et al.'s algorithm [2009] to compute them robustly and exactly, by replacing the metric from Euclidean to power distance.

3.3 Classification of Medial Spheres

For a medial sphere $\mathbf{m}_i = (\theta_i, r_i)$ that is tangent to the boundary surface at two or more points, its surface RPC $\partial \omega_i$ corresponds to an RPD constructed on the surface with respect to \mathbf{m}_i , which consists of a set of *Connected Components (CCs)* of surface regions, each of which contains one or more tangential surface points. For a smooth object, the medial spheres are organized into a small amount of classes [Giblin and Kimia 2004]. Since we are targeting for nonsmooth surfaces, we count the number of surface normals of those tangential points. We use our own notation $T_k^{n_1,\ldots,n_k}$ for the classification of a medial sphere \mathbf{m}_i . Here the subscript k represents the number of CCs of its RPC, and the superscript n_1,\ldots,n_k represents the number of surface normals for the tangential points in each CC.

For example, medial spheres of type $T_2^{1,1}$ lie on 2-manifold sheets, which are tangent to ∂S at exactly two distinct CCs and each CC contains only one tangential point (with its corresponding normal). Since this is the most ordinary case, to simplify the notations we remove all appearances of 1 from the superscript, *i.e*, $T_2^{1,1}$ will be simply denoted as T_2 . For most 3D shapes, the majority of medial spheres lies on type T_2 of *sheets*; the intersection of three or more local sheets forms a *seam* of T_3 spheres; and seams of type T_3 could



Fig. 4. The pipeline of our algorithm. Given a closed manifold 3D mesh surface (a) with sharp edges and corners pre-detected (marked in black lines), the initial medial spheres (b) are generated including zero-radius spheres (shown as red points) placed on convex sharp features and inner Voronoi balls (shown as blue spheres) of surface samples, as described in Sec. 4.1. Inner Voronoi balls protrudes from the surface are then updated and pushed to be tangential to the surface (c), as described in Sec. 4.2. An initial medial mesh (d) is then constructed based on RRT, and internal features (shown as red lines) are selected using a seam tracing algorithm, as described in Sec. 4.3. A refined medial mesh (d) with higher quality can be obtained from feature preservation strategies and a thinning process, as described in Sec. 4.4.

intersect at a *junction* sphere of type T_4 . Fig. 3 gives examples of these different classes.

In this paper we are handling non-smooth surfaces that could contain sharp edges and corners (either convex or concave). The tangential surface contact point on these sharp features does not have a unique and unambiguous normal. We define a convex sharp edge and concave sharp edge of input mesh ∂S as an edge subtending a dihedral angle less than $\pi - \phi$ and more than $\pi + \phi$ respectively [Abdelkader et al. 2020], where $\phi < \frac{\pi}{2}$ is an angle threshold used to bound the approximation error. Note that ϕ is a user-defined variable and users can also mark sharp features manually. A vertex of ∂S located on more than two sharp edges is defined as a *corner*.

- On a convex sharp feature, the medial spheres are of zero radii. The RPC of such zero-radius sphere have only one CC, but there could be two or three unambiguous normals associated with such feature point: sharp edges have two normals and sharp corners have at least three. Thus we denote such zero-radius medial spheres as T_1^2 (on a convex edge) and T_1^u ($u \ge 3$, on a corner). Note that a corner could be formed by a combination of both convex and concave sharp edges (Sec. 5.2 discusses these different cases and the mechanism to preserve the corner). For a corner that is formed by purely concave feature edges, there is no medial axis passing through it, thus the mechanism to handle the medial spheres adjacent to it is the same as the other concave features mentioned
- On a *concave sharp feature*, the medial spheres tangential to such features are not zero-radius. In fact, the normal direction of such tangential surface contacts are ambiguous (Sec. 4.1 gives detailed information of such medial spheres). Thus we use superscript c for annotating a CC that contains a tangential contact point on any concave external feature. See the examples of T_2^c and T_3^c in Fig. 3. Note that a medial sphere could be tangential to more than one concave feature edges on different CCs. Here we do not make distinct notations for them, as their sphere computation mechanism is similar to the case of single concave edge (see Sec. 4.2 for the details).

The CC containing an infinite number of tangential contact points will be annotated with superscript ∞ , e.g, for a cylinder shape, the CC of a medial sphere could contain the whole circular region that has an infinite number of tangential contact points with the sphere. In practice, since our input surface is represented as a triangle mesh, the number of tangential points for such a sphere is finite. Thus the superscript ∞ is only for notational purpose. See the examples of T_1^{∞} and T_2^{∞} in Fig. 3. They could be identified whenever a CC has more than one tangential point with diverse normals.

The external features of MAT include those convex sharp edges and their associated corners, and the internal features of MAT include those internal spheres located on seams (e.g., types of T_3 and T_3^c) as well as *junctions* (e.g., types of T_4 , T_4^c , and T_2^∞). Note that the examples of classes given in Fig. 3 are not meant to be complete. For example, there could be a corner denoted as T_1^6 formed by three convex edges and three concave edges, and there could be a junction sphere T_4^c that is tangential to two concave sharp edges and two other regular CCs, just to name a few. But notation-wise, all medial spheres can be represented with our CC-based classification. With this classification, the following sections discuss how we can compute these medial spheres based on such CC information, and how we can connect them to form a medial mesh while preserving both external and internal features of MAT.

THE COMPUTATIONAL PIPELINE

Our medial mesh computational pipeline consists of four major steps. Given a closed, manifold triangular mesh with sharp features pre-detected, the first step (Sec. 4.1) is to initialize the medial spheres. This includes placing sphere candidates on both non-feature and feature regions. The sphere candidates on non-feature regions are inner Voronoi balls generated with Delaunay Triangulation of surface samples. Special attention needs to be paid on initializing spheres tangential to concave sharp edges. The second step (Sec. 4.2) updates those inner sphere candidates as close as possible to their ground-truth positions and radii using our sphere updating strategy. Then in the third step (Sec. 4.3) we construct an initial medial mesh from these updated sphere candidates using the restricted regular triangulation (RRT) which is dual to RPD. Initial internal features could be detected through our seam tracing algorithm. In the fourth step (Sec. 4.4), the initial medial mesh is further refined through our internal feature preservation strategy and a thinning process to pursue the thinness property of medial mesh. Fig. 4 shows an illustration of our computational pipeline.

4.1 Medial Sphere Initialization

Non-feature Spheres: Our method starts with the generation of inner sphere candidates in those non-feature regions. A well-known technique is to use Voronoi balls (see Sec. 3.2) inside the shape ${\cal S}$ generated from some sampling points P on the boundary surface ∂S, and use the fast winding number [Barill et al. 2018] to keep only those inner Voronoi balls. Our assumption about the density of sample set P around smooth regions follows Amenta et al.'s work [2001a]. We use the local feature size function $LFS(\mathbf{x}): \partial S \to \mathbb{R}$ defined as the minimum Euclidean distance from a sample $x \in \partial S$ to the medial axis \mathcal{M} . A sample set P is an r-sample if any sample $\mathbf{x} \in P$ has a neighboring sample within its $rLFS(\mathbf{x})$ distance. Similar to existing Voronoi-based algorithms [Amenta et al. 2001a; Dey and Zhao 2002], we require $r \le 0.6$ for sufficiently approximating ∂S so that the generated Voronoi diagram can capture key information about shapes. Since the LFS around the convex sharp edges and corners are converging to zero, we stop the surface sampling that are within an η distance to these convex external features, where η is a distance threshold to avoid the density explosion to infinity around these external features.

Zero-radius Feature Spheres: To avoid dealing with infinity density around convex sharp features, we apply an adaptive medial re-sampling strategy. We first sample zero-radius medial spheres on those pre-detected convex sharp edges and corners, with the initial density same as its nearby surface samples that are η -distance away. After all those non-feature medial spheres are updated in the second step (Sec. 4.2), we might insert new zero-radius spheres to preserve the external feature. Specifically, for every non-feature medial sphere \mathbf{m}_x whose RPC is neighbor to those of zero-radius spheres on external features, we recursively check and add new zero-radius feature sphere if \mathbf{m}_x breaks the connection of external features. Please refer to Sec. 5.1 in more detail. Note that in this first step we only add zero-radius feature spheres using the initial density, while adaptive sphere insertion is recalled once the neighboring non-feature medial spheres are updated in the second step. Similarly, simply sampling zero-radius medial spheres on sharp corners is not enough for keeping the connectivity of medial axis. To preserve corner feature, we analyze the structure around corners and sample new medial spheres nearby to complete the medial structure. Please refer to Sec. 5.2 for detailed explanation.

Spheres Tangential to Concave Sharp Features: Even though concave sharp edges are not external features, a sharp change of surface normals around them leads to a smooth transition on medial axis (see Fig. 5 (b)), which requires dense samples of medial spheres around them. Since all concave sharp edges are pre-detected, we can sample dense medial spheres using sphere-shrinking algorithm [Ma et al. 2012]. We briefly introduce the algorithm as shown in Fig. 5 (a). For each pin-point ${\bf p}$ with corresponding normal ${\bf n}_{\bf p}$, a sphere ${\bf m}^0$ with large radius r^0 and center $\theta^0 = {\bf p} - r^0 {\bf n}_{\bf p}$ is initialized. This large sphere is then iteratively shrunk to approximate the medial sphere with two tangent points. For each iteration t+1, a new sphere ${\bf m}^{t+1}$ is found by performing a nearest point query from θ^t to the surface ∂S excluding point ${\bf p}$. The resulting nearest point ${\bf q}^{t+1}$ together with ${\bf p}$ is then used to compute the new sphere ${\bf m}^{t+1}$.

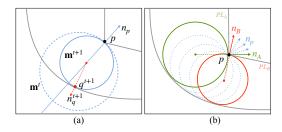


Fig. 5. Left: The *sphere-shrinking* algorithm [Ma et al. 2012] on concave regions in 2D. Right: Medial sphere initialization around concave region in 2D. For every pin-point \mathbf{p} sampled on concave sharp edge, we sample new normal \mathbf{n}_p between two normals \mathbf{n}_A and \mathbf{n}_B of planes adjacent to the concave edge, and initialize medial spheres using the sphere-shrinking algorithm. The denser the sampled pairs of $\{\mathbf{p},\mathbf{n}_p\}$, the smoother the final medial mesh around the concave edge we can get.

The iteration stops when the sphere \mathbf{m}^{t+1} is tangential on these two points and there is no other surface points closer to the sphere center. Our sampling strategy on concave sharp edges works as follows (see Fig. 5 (b) for a 2D illustration): for each pre-detected concave edge l, we sample pin-points \mathbf{p} densely on l and sample multiple normals $\mathbf{n}_{\mathbf{p}}$ in between two normals \mathbf{n}_{A} and \mathbf{n}_{B} of adjacent planes of the concave edge l. For each pair of $(\mathbf{p}, \mathbf{n}_{\mathbf{p}})$, we initialize a sphere that is large enough and then apply sphere-shrinking algorithm. The sphere-shrinking algorithm is guaranteed to converge [Ma et al. 2012] and it typically converges within a couple of iterations. It perfectly fits medial spheres with two tangential points, which are of types T_2 or T_2^c on sheets of the medial axis.

After the sphere initialization, however, inner Voronoi balls initialized for smooth regions are circumscribed over the surface sampling points in nature, making themselves often protrude the surface ∂S . Moreover, Voronoi-based generation in \mathbb{R}^3 is notorious for containing a large amount of Voronoi balls called *spikes* that are very close to the surface and far from the medial axis [Amenta et al. 2001a]. Even though medial spheres sampled on pre-detected concave edges are already at their ideal position using the sphere-shrinking algorithm, the internal features of MAT such as seams and junctions may not be fully represented. This is because the sphere-shrinking algorithm cannot sample spheres tangential to more than two surface points. Thus in the next step, these initial Voronoi spheres will be updated, and spikes will be removed.

4.2 Medial Sphere Update

Given the initial medial spheres estimated from inner Voronoi balls, we apply local operations to update their positions and radii as close as possible to the ground truth. Ideally we expect all medial spheres to be tangent to at least two points on surface ∂S without protrusion.

For inner Voronoi balls initialized with two tangent points, *i.e.*, type T_2 or T_2^c on medial sheets, we use sphere-shrinking algorithm [Ma et al. 2012] as described above to update them. Although this algorithm is highly efficient, it cannot handle spheres with more than two tangent points, *i.e.*, T_3 and T_3^c on medial seams, or T_4 and T_4^c on medial junctions. The problem is more challenging than the case of two tangent points, because it is very difficult to find an exact

pin-point on either the seam or junction to determine the sphere. In this paper we formulate the updating algorithm for spheres with any number of tangent points as a continuous optimization problem, and name it as multi-tangent sphere optimization algorithm.

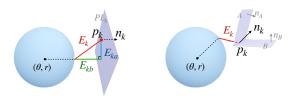


Fig. 6. Illustration of energy terms for our multi-tangent sphere optimization. Left: Energy term \mathcal{E}_k as the sum of \mathcal{E}_{ka} and \mathcal{E}_{kb} when the tangential point \mathbf{p}_k is on a plane \mathbf{PL}_k . Right: energy term E_k when the tangential point \mathbf{p}_k is on a concave sharp edge \mathbf{l}_k .

Our multi-tangent sphere optimization is inspired by the property of surface RPC. Each CC of sphere m defines at least one tangent plane or tangent concave edge of m with their tangential points \mathbf{p}_k and normals \mathbf{n}_k , where k is the index of all tangential points belonging to m. In this problem, not only the sphere $\mathbf{m} = (\theta, r)$ needs to be solved, but also the exact positions and normals (\mathbf{p}_k , \mathbf{n}_k) of all tangential points are unknown. We define the following quadratic energy for our multi-tangent sphere optimization:

$$E(\theta, r, \{\mathbf{p}_k, \mathbf{n}_k\}_{k=1}^N) = \sum_{k=1}^N E_k,$$

$$E_k = ||\theta + r\mathbf{n}_k - \mathbf{p}_k||^2.$$
(5)

Here N is the total number of tangential points for this sphere, and E_k is the energy term defined using the tangent pair $(\mathbf{p}_k, \mathbf{n}_k)$. Note that tangential point \mathbf{p}_k can be either on a plane (Fig. 6 (a)) or on a concave sharp edge (Fig. 6 (b)). We define E_k as the squared Euclidean distance from \mathbf{p}_k to the expected point of tangency on sphere. This energy is zero when the sphere is exactly tangential to the surface at point \mathbf{p}_k .

As shown in Fig. 6 (a), we can split the energy term E_k as the sum of two sub-terms E_{ka} and E_{kb} :

$$E_{kb} = ||(\mathbf{p}_k - \boldsymbol{\theta})^{\top} \mathbf{n}_k - r||^2,$$

$$E_{ka} = E_k - E_{kb}.$$
(6)

Intuitively E_{kb} is the squared distance from sphere **m** to plane PL_k along the direction of normal n_k , and E_{ka} is the squared distance on the plane. We modify the energy in Eq. (5) by incorporating two weights λ_a and λ_b to balance the significance of these two sub-terms:

$$E(\theta, r, \{\mathbf{p}_k, \mathbf{n}_k\}_{k=1}^N) = \sum_{k=1}^N \lambda_a E_{ka} + \lambda_b E_{kb}.$$
 (7)

Inspired by the classic iterative closest point (ICP) algorithm [Chen and Medioni 1992], giving a higher weight λ_b to the tangential squared distance E_{kb} could potentially speed up the convergence of optimization. We use λ_a = 0.01 and λ_b = 1 in all of our experiments.

If \mathbf{p}_k is on a smooth surface region (Fig. 6 (a)), then its normal \mathbf{n}_k is fully determined by its position. If \mathbf{p}_k is on a concave sharp edge l_k (Fig. 6 (b)), then n_k can be any direction in between two normals \mathbf{n}_A and \mathbf{n}_B of adjacent planes of the concave edge. These two normals give us a bound of possible directions that the medial sphere can be tangent to.

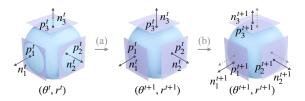


Fig. 7. Iterative two-stage optimization for medial spheres. (a) The sphereupdating stage which locks the aggregated tangent pairs $(\mathbf{p}_k, \mathbf{n}_k)$, k = 1...Nand update the medial sphere (θ, r) ; (b) The tangent-updating stage that fix the previously updated medial sphere then update each tangent pair. Both stages optimize the same energy function in Eq. (7).

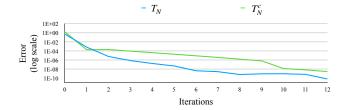


Fig. 8. Convergence of the sphere updating optimization of Eq. (7) through iterations (in log scale for better illustration). The error is measured as the squared distance from the sphere to the surface. We show the average errors collected from all models shown in Fig. 17 aggregated in two sphere types: T_N and T_N^c , where N > 2. Note that T_2 and T_2^c are not collected since we use sphere-shrinking algorithm [Ma et al. 2012] in Sec. 4.1

Since both the sphere $\mathbf{m} = (\theta, r)$ and the tangent pairs $(\mathbf{p}_k, \mathbf{n}_k)$, k = 1...N are to be determined, we design our multi-tangent sphere optimization algorithm as an iterative two-stage optimization process. For each new iteration t + 1, we decompose the optimization into a sphere-updating stage and a tangent-updating stage:

- During the sphere-updating stage (Fig. 7 (a)), we fix all tangent pairs $(\mathbf{p}_k^t, \mathbf{n}_k^t)$, k = 1...N from the previous iteration t, and update the sphere (θ^{t+1}, r^{t+1}) by minimizing the energy in Eq. (7). Since the energy is a quadratic function of θ^{t+1} and r^{t+1} , they can be solved by a simple linear equation.
- During the tangent-updating stage (Fig. 7 (b)), we lock the sphere (θ^{t+1}, r^{t+1}) from the first stage and update all tangent pairs $(\mathbf{p}_k^{t+1}, \mathbf{n}_k^{t+1})$, k = 1...N. If \mathbf{p}_k^t is on a plane \mathbf{PL}_k of the surface, we search the $\operatorname{\mathcal{K}\text{-}ring}$ neighboring surface triangles around \mathbf{p}_{k}^{t} ($\mathcal{K}=2$ in our experiments), and find a point within these triangles that minimizes the quadratic energy E as the new tangent point. As each triangle has a fixed normal, searching inside a triangle is equivalent to solving a linear system of barycentric coordinates. If \mathbf{p}_k^t is on a concave sharp edge \mathbf{l}_k , the optimal tangent point \mathbf{p}_k^{t+1} is simply the projection of the sphere center θ^{t+1} onto the concave sharp edge, and the optimal normal \mathbf{n}_k^{t+1} is the normalized direction

of $\mathbf{p}_k^{t+1} - \boldsymbol{\theta}^{t+1}$. Note that \mathbf{n}_k^{t+1} will be clamped to boundary normal once out of range of \mathbf{n}_A and \mathbf{n}_B .

Since both stages decrease the quadratic energy E in Eq. (7), the optimization is converging through iterations. We visualize the convergence of our iterative optimization scheme in Fig. 8. Note that the iteration process does not require a re-calculation of RPD. All inner Voronoi balls will be updated iteratively until their energy functions E are smaller than a threshold ϵ (10⁻⁴ in our experiment). A converging energy larger than ϵ indicates that such a sphere does not exist to be tangential to the set of given planes and/or concave edges, and it will be removed consequently. This energy-checking will become very useful in our sphere-insertion operations, e.g., for external corner preservation described in Sec. 5.2.

It is worth noting that spheres of type T_1 are normally *spikes*, because they only touch one side of the surface and has only one tangent point and normal (see the right inset figure). Even though our method is not specifically designed for pruning spikes, a by-product of our computed RPC is to robustly detect this type of medial spheres and remove them.



4.3 Medial Mesh Initialization

After the sphere update, there are some non-feature spheres that break the connectivity of nearby external features (details in Sec. 5.1). We will re-sample zero-radius feature spheres in those cases. Once all spheres (including zero-radius spheres and update inner Voronoi balls) are at their ideal positions (Fig. 4 (c)), we build their connectivity to form a structured medial mesh. Our topological connection method is inspired by the *power shape* [Amenta et al. 2001b] which is a subset of the *regular triangulation* (RT) dual to the *power diagram* (PD). To construct the medial mesh \mathcal{M}_s , we compute the *restricted regular triangulation* (RRT), by selecting a subset of simplices in RT, whose dual elements in PD have non-empty intersections with the input shape \mathcal{S} . Specifically, we check the RT simplicies in the order of tetrahedra, triangles, and edges:

- The dual of an RT tetrahedron is a vertex in PD. If this dual vertex is inside the shape S, then we keep this tetrahedron together with all of its triangles and edges in the medial mesh.
- If the dual vertex of the RT tetrahedron is outside the shape, then we check all of its four triangles. The dual of an RT triangle is an edge segment in PD. If this dual edge segment has any intersection with the shape S, then we keep this triangle together with all of its edges in the medial mesh.
- If the dual edge segment of the RT triangle is outside the shape, then we check all of its three edges. The dual of an RT edge is a polygonal face in PD. If there is any intersection between the dual polygonal face with the shape *S*, then we keep this RT edge in the medial mesh.

Note that in this process we only need to compute RT as dual of PD and compute its restriction to the shape $\mathcal S$ by checking the intersection between its dual edge segment or polygon with input surface ∂S . In this way we do not need to compute the volumetric RPD which requires a tetrahedralization of the input shape $\mathcal S$ and the cutting of those tetrahedra with PD. It is worth mentioning that,

similar to power shape [Amenta et al. 2001b], our initial medial mesh \mathcal{M}_s generally contains some flat but solid tetrahedra. However, the medial axis of a three-dimensional shape \mathcal{S} should be a collection of two-dimensional sheets, *i.e.*, they should be *thin* without any solids. All these tetrahedra in the initial medial mesh will be pruned in our thinning process in the fourth step (Sec. 4.4).

Once the initial medial mesh is constructed, we can trace the internal features using a seam tracing algorithm similar to Culver et al.'s method [2004]. We first classify initial medial spheres based on their CCs and tangent points, and identify those medial spheres on seams (*i.e.*, type T_3 or T_3^c). Our seam tracing algorithm starts from any seam sphere and expand to its adjacent feature spheres. The detailed seam tracing algorithm is provided as Alg. 1 in Supplementary Material. In this way we can detect an initial set of internal features (Fig. 4 (c) red lines) which should be further refined in the next step.

4.4 Medial Mesh Refinement

The initialized and updated spheres in the first two steps cannot guarantee there are sufficient spheres sampled on the internal features, such as seams and junctions. This deficiency usually happens in local regions where an ill-posed connection is caused by two medial spheres that lie on two different medial sheets (see Fig. 9). We first detect these connections under the help of their RPCs, then sample new internal feature spheres using the tangential surface contact points aggregated from two ill-connected spheres. Sec. 5.3 explains this refinement in detail.

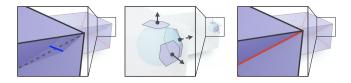


Fig. 9. Left: Ill-posed connection (blue line) of two T_2 spheres on two different medial sheets. Middle: A new feature sphere of type T_3 is inserted to preserve the internal feature. Right: The internal feature (red line) preserved after our refinement strategy.

The initial medial mesh still contains some flat but solid tetrahedrons as mentioned above. We adapt the thinning algorithm proposed by Liu et al. [2010] that prunes simple pairs of simplices in the medial mesh. A simple pair (x, y) is a pair of simplices such that y is on the boundary of *x* and there is no other cell in the complex with y on its boundary. Fig. 11 (a) shows a single tetrahedron t in medial mesh as an example, where the tet-face pair (t, f) is a simple pair, but the face-edge pair (f, e) is not. It is shown by Ju et al. [2007] that removing a simple pair does not change the topology even when multiple simple pairs are removed together. This thinning is a pure topological operation that can help us remove tetrahedra from the medial mesh. However, to remove a simple tet-face pair, there could be multiple potential choices. The different order of such removal operation will result in different geometry of the medial mesh, even though their topologies are equivalent. When deciding which simple pair to remove given multiple choices, we shall introduce a quantitative measure for ranking those tetrahedral faces, so

that the faces that are less important are prioritized over others to be removed first.



Fig. 10. Illustration of restricted power segment (RPS). (a) Given a medial face of three medial spheres \mathbf{m}_i , \mathbf{m}_j and \mathbf{m}_w , and (b) their restricted power cells (RPCs), its RPS is shown as the green dotted segment γ_{ijw} .

Consider a medial face f_{ijw} connecting three medial spheres \mathbf{m}_i , \mathbf{m}_{i} and \mathbf{m}_{w} (see Fig. 10). Since f_{ijw} is a triangle of RT, its dual Γ_{ijw} is a line segment of PD. We define the intersection of Γ_{ijw} with the shape S as the restricted power segment (RPS) γ_{ijw} :

$$\gamma_{ijw} = \Gamma_{ijw} \cap \mathcal{S}. \tag{8}$$

It is worth mentioning that the endpoints of an RPS can be either on the surface ∂S or inside S. This is because its endpoint could be a dual vertex of an RT tetrahedron, which could be potentially inside S (as mentioned in Sec. 4.3).

For medial face f_{ijw} , its dual RPS γ_{ijw} must be perpendicular to the triangle face created by three medial sphere centers (θ_i , θ_j , θ_w). For a local region of medial mesh that is "thin", the length of γ_{ijw} approximates the local thickness of shape, see Fig. 10 the segment in green. Therefore, we defines the importance factor α_{ijw} of a given medial face f_{ijw} as the ratio of the length of γ_{ijw} over the average diameter of three medial spheres.

For a local medial mesh region that has a flat tetrahedron, every triangle f_{ijw} of the tetrahedron has a dual RPS. We rank these medial triangles by their importance factors α_{ijw} in ascending order. Fig. 11 (b) shows a 2D example where there is a medial triangle (instead of a tetrahedron in 3D) in the neighborhood of a concave sharp feature. The 2D medial triangle $(\mathbf{m}_A, \mathbf{m}_B, \mathbf{m}_C)$ exists when the undesirable edge (θ_A, θ_C) exists, which means its dual RPS γ_{AC} exists. In fact, γ_{AC} degenerates to a point p near the concave sharp feature, making its importance factor α_{AC} to be close to zero. The reason for this degeneracy is because all three spheres pass through the concave sharp feature, making the two RPC of spheres \mathbf{m}_A and \mathbf{m}_C to be adjacent at the concave point p. The other two RPS γ_{AB} and γ_{BC} are shown as green and red dotted segments, respectively. Their importance factors α_{AB} and α_{BC} are both close to one. Apparently medial edges $\theta_A \theta_B$ and $\theta_B \theta_C$ are more important than edge $\theta_A \theta_C$ after ranking their dual RPS with their importance factors.

Our geometry-guided thinning algorithm starts from raking all triangles of any tetrahedron in the medial mesh using their importance factors. Then we remove tet-face simple pairs with least importance in each iteration until all tetrahedra are pruned. After removing all tet-face simple pairs, we continue pruning those faceedge simple pairs that belong to the original tetrahedra. To avoid over-pruning for models whose medial mesh boundaries are not external features, we use a target important factor σ as a stop sign

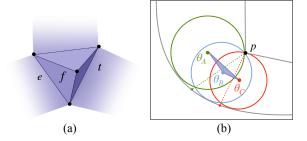


Fig. 11. Illustration of some tetrahedron in 3D medial mesh as a 3dimensional cell of (a), and in 2D as a 2-dimensional cell of (b).

so that the face-edge pair will be not deleted when the importance factor of the current face is beyond this target. Note that σ will only impact the pruning of face-edge pairs but not tet-face pairs, the choice of σ is discussed in Sec. 6.4. This pruning operation results in a 3D medial mesh that is "thin" with no three dimensional cells, and at the same time, maintains high-quality geometry of the medial mesh. We provide the detailed algorithm as Alg. 2 in Supplementary Material, and an ablation study result in Sec. 6.4.

5 FEATURE PRESERVATION

The medial feature spheres sampled during initialization (Sec. 4.1) are not guaranteed to preserve external features after the update of non-feature medial spheres (Sec. 4.2). Also the initial internal features traced in the initial medial mesh (Sec. 4.3) require further refinement (Sec. 4.4). The connection of two neighboring feature spheres are likely to be destroyed by some non-feature spheres nearby, which results in a fracture of medial features, either external or internal. This is caused mainly by one reason: the feature spheres are not sampled sufficiently in a local region so that the RPC of non-feature spheres may "invade" the neighboring RPCs of two feature spheres. Our solution for this issue is to detect the local regions which lack feature spheres, and then add new spheres on them. We will discuss the details of our preservation strategy for both external and internal features.

5.1 Preserving External Edge Features

During the generation of initial medial mesh, we sample zero-radius spheres on convex sharp edges to avoid handling infinite sampling density on the input surface. The local external feature could be destroyed if an RPC of any non-feature sphere "invades" the RPCs of two neighboring feature spheres on sharp edge (see Fig. 12 left). This is due to the deficiency of feature spheres on external edge features. Therefore, our adaptive re-sampling strategy for maintaining the external features of convex sharp edges is to detect such cases of non-feature spheres then recursively adding new feature spheres until all external edge features are preserved.

We observe that the relationship between RPCs of three spheres can help us find those non-feature spheres that destroy the connectivity of two neighboring zero-radius medial spheres on convex sharp edges. See Fig. 12 for illustrations. For any non-feature medial sphere $\mathbf{m}_{x} = (\boldsymbol{\theta}_{x}, r_{x})$ whose RPC is neighboring to the RPCs of two zero-radius medial spheres $\mathbf{m}_a = (\theta_a, 0)$ and $\mathbf{m}_b = (\theta_b, 0)$, where

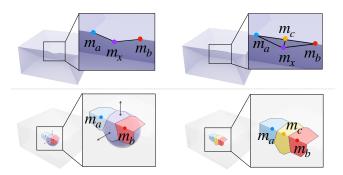


Fig. 12. Illustration of how the RPC of a non-feature sphere \mathbf{m}_x "invades" the RPCs of two neighboring feature spheres \mathbf{m}_a and \mathbf{m}_b on a sharp edge (bottom left), which could destroy the connectivity of $\mathbf{m}_a\mathbf{m}_b$ in the medial mesh (top left). We should add a new feature sphere \mathbf{m}_c (right) and then recursively check if the sampling is enough.

 $\theta_a \theta_b$ is a feature edge to be preserved. The medial sphere \mathbf{m}_x breaks the connectivity of $\theta_a \theta_b$ if the following inequation is true:

$$\theta_x^{\top} \theta_x - (\theta_a + \theta_b)^{\top} \theta_x + \theta_a^{\top} \theta_b \le r_x^2. \tag{9}$$

The detailed proof is given in the Supplementary Material Sec. 2.

For every non-feature sphere \mathbf{m}_x , we search for its \mathcal{K} -nearest external edge features $\mathbf{m}_{a_k}\mathbf{m}_{b_k}$ ($k=1...\mathcal{K}$) and check if \mathbf{m}_x destroy the connectivity of them. If it breaks, we add a new feature sphere \mathbf{m}_{c_k} in the middle of edge $\mathbf{m}_{a_k}\mathbf{m}_{b_k}$ and recursively check edges $\mathbf{m}_{a_k}\mathbf{m}_{c_k}$



and $\mathbf{m}_{b_k}\mathbf{m}_{c_k}$. This recursion will be infinite only when \mathbf{m}_x protrudes from the convex sharp edge (see the right inset figure) and any new feature sphere \mathbf{m}_{c_k} could be contained inside \mathbf{m}_x . Therefore in this case the medial sphere \mathbf{m}_x should be removed. However, this is unlikely to happen since all non-feature spheres have been updated to their ideal position in the second step, so they would not protrude from the convex sharp edges.

5.2 Preserving External Corner Features

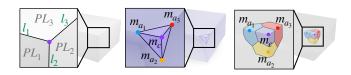


Fig. 13. Left: A convex corner formed by three convex edges \mathbf{l}_i and three incident planes \mathbf{PL}_i (i=1,2,3). Middle: The structure of the "corner cap" where any two neighboring external feature spheres \mathbf{m}_{a_i} and \mathbf{m}_{a_j} are connected, created a medial face of $\triangle(\mathbf{m}_c\mathbf{m}_{a_i}\mathbf{m}_{a_j})$. Right: The surface RPC of these four external feature spheres.

An external corner is typically formed by three or more sharp edges, including either convex edges or concave edges. As shown in Fig. 13, zero-radius medial spheres sampled on one sharp convex edge may get entangled with the ones on neighboring sharp convex edges, which may result in a "corner cap" in the resulting medial

mesh. Since each sharp edge \mathbf{l}_k is formed by two adjacent planes \mathbf{PL}_{k_1} and \mathbf{PL}_{k_2} , it must define a sheet of medial axis passing through this sharp edge, with medial spheres tangent to its two adjacent planes. Three sheets potentially intersect and form a seam. For example in Fig. 14 I, any convex edge \mathbf{l}_k forms a sheet s_k on medial axis, and three sheets join at a seam e_{123} .

One possible solution to remove the "cap" around corners is to directly sample new medial spheres on potential seams using our internal feature preservation strategy described in Sec. 5.3. However, such addition of new spheres on seams might break the connectivity of sharp edges nearby (Sec. 5.1), which would require adding more samples on sharp edges in order to preserve them. This would cause an infinite loop and the sphere sampling density could run into infinity, as two sharp edges are approaching closer to a corner. To solve this challenge, we select a small region within distance δ from the given corner \mathbf{m}_c , and propose a corner preservation strategy to approximate medial mesh structures in this small region around the corner. Our corner preservation scheme works by recursively tracing sheets of medial axis, starting from those convex sharp edges, until we found their intersecting seams. It consists of the following three steps.

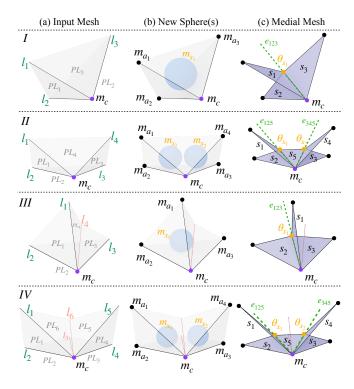


Fig. 14. Some examples of corners incident to convex and/or concave sharp edges. For each example, we show their valid seams represented as tangent collections symbolically. Row $I: e_{123} = (PL_1, PL_2, PL_3)$. Row $II: (1) e_{125} = (PL_4, PL_1, PL_2)$; (2) $e_{345} = (PL_2, PL_3, PL_4)$. Row $III: e_{123} = (PL_2, I_4, PL_1)$. Row $IV: (1) e_{125} = (PL_6, PL_1, PL_2)$; (2) $e_{345} = (PL_3, PL_4, PL_5)$.

5.2.1 Enumerating Initial Set of Sheets. We first sort neighboring sharp edges \mathbf{l}_k (k=1...N), both convex and concave, together with their adjacent planes \mathbf{PL}_k in a counter-clockwise order. In this

way we can enumerate an initial set of potential sheets using the tangent planes incident to all convex sharp edges. Note that concave sharp edges will not contribute any new sheet as they only create smooth transitioning conic sections on existing sheets, and spheres tangential to only one concave edge are of type T_1^c as spikes.

Take Fig. 14 III as an example. Corner \mathbf{m}_c is where three convex edges (l₁, l₂ and l₃ in green) and one concave edge (l₄ in red) converge. Also \mathbf{m}_c has four incident planes PL_k , $k \in \{1...4\}$, where PL_3 and PL4 are shared by a concave edge l4. There are three possible sheets s_k traced from convex edges l_k , k = 1, 2, 3. For example, sheet s₁ is formed by a set of medial spheres tangential to two tangent planes PL4 and PL1 adjacent to convex edge l1. So we can have symbolic representations for these sheets: $s_1 = (PL_4, PL_1)$, $s_2 = (PL_1, PL_2)$ PL_2), and $s_3 = (PL_2, PL_3)$.

- 5.2.2 Finding Seams and New Sheets Recursively. If two neighboring sheets in counter-clockwise order intersect, we will be able to find a medial sphere on the intersecting seam through their aggregated set of tangent planes and/or concave edges. In addition, we can find out another new sheet that intersects with them on the same seam. In the example of Fig. 14 III, we can find their potential intersecting seams as follows:
 - (1) Checking potential intersection between s_1 and s_2 :
 - (a) Using tangent collection (PL₄, PL₁, PL₂) or
 - (b) Using tangent collection (l₄, PL₁, PL₂);
 - (2) Checking potential intersection between s_3 and s_1 :
 - (c) Using tangent collection (PL1, PL2, PL3) or
 - (d) Using tangent collection (PL₁, PL₂, l₄).

Here (b) and (d) defines the same seam symbolically so we can merge them as e_{123} . Each tangent collection of three elements (tangent planes and/or concave edges) imply a potential medial sphere on the seam of medial axis. We can obtain the sphere \mathbf{m}_{x_i} by minimizing the energy of Eq. (7) as described in Sec. 4.2. If the optimized energy is larger than threshold ϵ , it means there does not exist a sphere that can be tangential to these tangent collections. The example in III only have one valid seam sphere based on collections (b) and (d), while the other two collections (a) and (c) cannot lead to any medial sphere based on their optimized energies. Once a seam sphere is computed on collection (b), based on the tangent information we can obtain a new sheet (l_4, PL_2) that is intersecting with both s_1 and s_2 on the same seam. In this case the new sheet is exactly $s_3 = (PL_2, PL_3)$ because the concave edge l4 and tangent plane PL3 are adjacent, so it will not be included further.

This process is implemented as a circular queue of potential sheets ordered counter-clockwise. Every time we pop out a sheet, we check its intersection with the next sheet in the queue. If two of them forms a seam, we sample a new medial sphere with two sheets' info recorded, and push the new sheet candidate (if unique) into the queue. Otherwise, we push it back into the queue and pop out the next sheet to check. This process continues until the queue is empty.

5.2.3 Connect the Spheres on Sheets and Seams to form a Medial Mesh around the Corner. Once all seams are found with their medial spheres sampled, our third step is to construct an approximated medial mesh around the corner. For each convex sharp edge, we sample a zero-radius medial sphere at distance δ from the corner. For example, in Fig. 14 IV we sample \mathbf{m}_{a_1} , \mathbf{m}_{a_2} , \mathbf{m}_{a_3} , and \mathbf{m}_{a_4} on convex edges l1, l2, l4, and l5, respectively. For each sheet found in the above process, we connect its two medial spheres with the corner to form a medial triangle, e.g., , $\mathbf{m}_{x_1}, \, \mathbf{m}_{x_2}$ and the corner \mathbf{m}_c form a triangle for the sheet s_5 ; \mathbf{m}_{x_1} , \mathbf{m}_{a_1} and the corner \mathbf{m}_c form a triangle for the sheet s_1 , etc.

In this way all sheets around the corner can be constructed as medial triangles, with all seams represented as edges between neighboring triangles. In Sec. 6.5 we show our experiment results on shapes with different corner cases. Note that a corner could be just the tip of a cone shape without any neighboring sharp edges (see Fig. 22 (e)), in which case our medial mesh computed with RRT in Sec. 4.3 can already preserve it as long as there is a zero-radius medial sphere sampled on the tip.

5.3 Preserving Internal Features

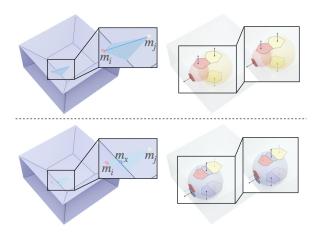


Fig. 15. The local region near two medial spheres m_i and m_j that are on two different medial sheets. There is a lack of feature spheres in this local region, causing m_i and m_j to be connected (top). A new feature sphere m_x is added to preserve the internal feature of medial axis (bottom).

Internal feature preservation requires sufficient sampling of internal feature spheres. Therefore, our initial internal features detected from the initial medial mesh need to be refined after detecting the regions that lack feature spheres. We have observed that the deficiency of internal feature spheres normally occurs when two connected non-feature medial spheres belong to different medial sheets. Based on this observation, we propose an internal sphere insertion strategy that is able to directly insert medial spheres without relying on additional surface samples (for computing Voronoi balls).

Suppose two medial spheres \mathbf{m}_i and \mathbf{m}_j are connected in the medial mesh. We are expecting every CC of \mathbf{m}_i must be adjacent to one corresponding CC of \mathbf{m}_i , which means sphere \mathbf{m}_i and \mathbf{m}_i are on the same medial sheet. Otherwise they should not be connected since they are on different medial sheets. A deficiency of feature spheres is detected in this local region if m_i and m_j belong to two different medial sheets on the medial mesh. In this case, a new feature sphere \mathbf{m}_x will be inserted using the aggregated connected components of \mathbf{m}_i and \mathbf{m}_i . Fig. 15 shows an illustration of such case. Two T_2 medial spheres \mathbf{m}_i and \mathbf{m}_i are connected so that the internal feature (black solid line) cannot be preserved as expected. Only one CC of \mathbf{m}_i is adjacent to one CC of \mathbf{m}_j , and their another CCs are on different surface regions and not adjacent to each other (see Fig. 15 right column). This indicates that \mathbf{m}_i and \mathbf{m}_j are on two different medial sheets, and there is a deficiency of feature spheres in this local region. Using the aggregated CCs of these two medial spheres, we can insert a new medial sphere \mathbf{m}_x of type T_3 that has tangential contacts with the surface at three points using the same method in Sec. 4.2.

6 EXPERIMENTS

We implement our algorithm in C++, using CGAL for triangulation calculation, and Eigen for linear algebra routines. The exact calculation of restricted power diagrams extends the Voronoi package in Geogram¹. We run our experiments on a computer with a 3.60GHz Intel(R) Core(TM) i7-9700K CPU and 32 GB memory. All models used in this paper are from the ABC dataset [Koch et al. 2019] and their sizes are normalized to the [0, 10] range.

Evaluation Metrics. We use the two-sided Hausdorff distance error, denoted as ϵ , to assess the surface reconstruction accuracy using the generated medial meshes. ϵ^1 is the one-sided Hausdorff distance from the original surface to the surface reconstructed from MAT, and ϵ^2 is the distance in reverse side. We also directly evaluate the difference between the approximated medial mesh and the groundtruth medial axis for some input surfaces. The ϵ_{ma}^1 is the one-sided Hausdorff distance from the ground-truth medial axis to the approximated medial mesh, and ϵ_{ma}^2 is the distance in reverse side. Note that we manually generate the ground-truth medial axis for some simple shapes, such as those in Fig. 2, Fig. 16, and Fig. 21. For sampling-based methods, we also show #v as the number of surface samples used. We show #s as the number of medial spheres for the medial meshes generated from each method. All Hausdorff distances are evaluated as percentages of the distances over the diagonal lengths of the models' bounding boxes.

We show a set of qualitative results on various 3D CAD models in Fig. 17 with their running time statistics summarized in Table 1. More detailed views of those computed medial mesh and their extracted medial features are shown in the supplementary video. The shape reconstruction errors of those models listed in Fig. 17, measured by the Hausdorff distance error, are given in Table 2.

6.1 Comparison with the CAD_MEDIAL Method

To our best knowledge, CAD_MEDIAL [Dey et al. 2003] is the only method that works on preserving external features of medial axis so far. However, their sampling condition near sharp edges is extremely strict and very hard to be achieved if the sharp edge is not a straight line. Fig. 16 shows the visual comparison of the medial mesh quality of ours and CAD_MEDIAL with different numbers of surface samples used on a CAD model. It also shows the Hausdorff distances w.r.t. its ground truth medial axis. It can be seen that increasing the number of surface samples for CAD_MEDIAL would make their generated medial mesh more complete, however, the sampling density

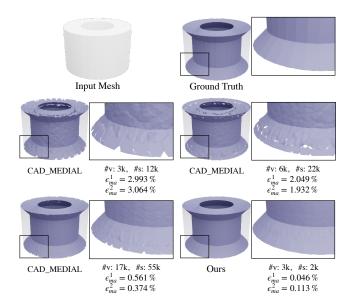


Fig. 16. Comparison with CAD_MEDIAL [Dey et al. 2003] on a CAD model with ground truth medial axis.

Table 1. Statistics of our running time in seconds. $\#v^*$ is the number of vertices of original model, #v is the number of surface samples used to generate initial Voronoi balls, #f is the number of triangle faces in the input mesh, #s is the number of generated medial spheres, and #t is the number of active tetrahedrons before thinning process in Sec. 4.4. Note that the running time of calculating surface RPD relates to #v, #f, and #s using the clipping algorithm [Yan et al. 2009]. And the running time of thinning process in S4 relates to #t. S1 is the running time of calculating initial medial sphere centers (Sec. 4.1). S2 is the running time of updating medial spheres (Sec. 4.2). S3 is the running time of calculating initial medial mesh (Sec. 4.3). S4 the running times for refining medial mesh (Sec. 4.4). The model's ID#t corresponds to those shown in Fig. 1 and Fig. 17.

Model	#v*	#v	#f	#\$	#t	S1	S2	S3	S4	Total
020	21k	5k	41k	11k	4k	2.8	1.4	4.9	6.8	15.9
068	7k	2k	14k	5k	773	0.8	0.3	1.2	4.6	8.9
077	19k	5k	38k	11k	17k	2.5	2.1	4.7	50.3	59.6
125	5k	5k	10k	6k	4k	0.7	0.5	1.1	7.6	9.9
128	25k	5k	51k	31	28k	3.8	9.8	12.6	84.5	110.7
129	21k	4k	42k	18k	11k	2.8	2.6	6.4	16.5	23.1
152	41k	9k	82k	39k	37k	6.3	4.8	27.1	118.6	156.8
168	22k	4k	44k	11k	9k	2.8	1.0	5.5	14.5	23.8
287	22k	5k	45k	12k	12k	2.9	1.7	6.5	23.2	34.3
329	5k	2k	9k	6k	507	0.7	0.3	1.5	4.2	6.7
801	17k	4k	34k	9k	4k	2.0	0.9	4.3	6.3	13.5
802	11k	3k	21k	6k	5k	1.3	0.7	2.5	6.7	11.2

around sharp edges needs fine tuning and very likely to generate incomplete structures. In contrast, our method does not require strict sampling condition around those non-smooth surface regions and produces more accurate approximation of convex external features with smaller Hausdorff distances.

6.2 Comparison with Sampling-based Methods

We compare our method with Power Crust (PC) [Amenta et al. 2001a] regarding the generated medial mesh (Fig. 21), the surface

 $^{^{1}} Geogram: http://alice.loria.fr/software/geogram/doc/html/index.html \\$

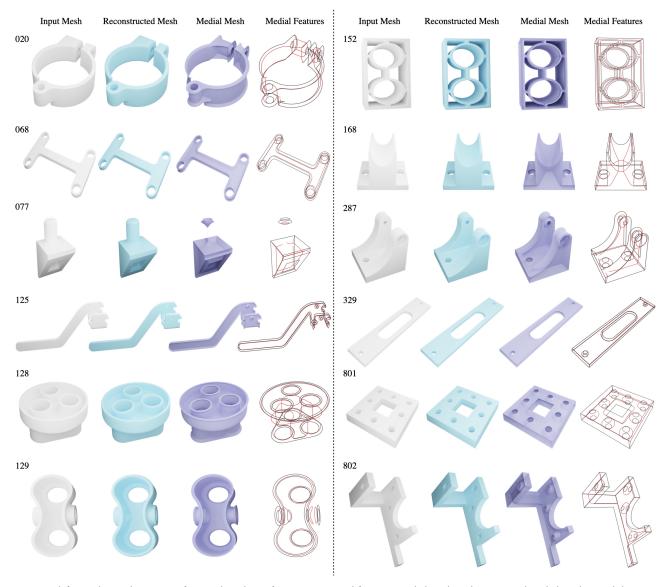


Fig. 17. From left to right are the input surface meshes, the surfaces reconstructed from our medial meshes, the generated medial meshes, and the extracted medial features. For the medial features, the black curves are the external features and the red curves are the internal features. For reconstructed surfaces, their reconstruction errors are given in Table 2.

reconstruction from MAT (Fig. 18), and the feature preservation quality (Fig. 19). Similar to other sampling-based methods, the quality of medial mesh generated using PC would improve when the surface sampling density increases. However, their method cannot preserve any medial feature and the generated medial mesh is not thin with large number of flat tetrahedrons.

We also experiment with the SAT method [Miklos et al. 2010] using two values of the sampling distance parameter: δ = 0.03 and δ = 0.01, and setting the scale parameter to s = 1.0. The smaller δ yields a good reconstruction precision. However, it favors a dense representation with a large number of medial spheres. The qualitative and quantitative comparison results are shown in Fig. 18. SAT cannot preserve external features of input mesh surface in their MA results as the medial structure is not complete around convex sharp edges of input surfaces. Even though SAT generates promising result when preserving internal features on some models, it requires large amount of medial spheres (i.e., 256k and 127k in Fig. 19 with $\delta = 0.01$) comparing to our method (i.e., 15k and 8k in Fig. 19). When the number of medial spheres is not adequate even with same sampling parameter (i.e., 7k and 57k in Fig. 18 with δ = 0.01), SAT generates ill-posed faces around internal features. In addition, SAT also routinely produce topological errors, as shown by the Euler characteristic in Fig. 19.

Table 2. Quantitative comparison on shape reconstruction error among PC [Amenta et al. 2001a], SAT [Miklos et al. 2010], VC [Yan et al. 2018] and ours. #s is the number of generated medial spheres, ϵ is the two-sided Hausdorff distance between original surface and reconstruction (maximum of ϵ^1 and ϵ^2 described in Sec. 6), and E is the Euler characteristic. We also show ground truth Euler characteristic as "GT E". For CAD models, our method always gives the best reconstruction results with correct Euler characteristics, and with smaller amount of medial spheres generated.

Model	PC			SAT			VC			Ours		
(GT E)	#s	ϵ	E	#s	ϵ	E	#s	ϵ	E	#s	ϵ	E
020 (-1)	35k	0.733	78k	137k	0.715	97	34k	1.319	-1	11k	0.400	-1
068 (-3)	11k	0.577	20k	85k	0.572	44	77k	0.722	-3	5k	0.550	-3
077 (1)	32k	2.508	86k	235k	1.492	140	72k	2.822	1	11k	1.381	1
125 (-2)	7k	0.708	8k	51k	0.406	31	10k	0.501	-2	6k	0.198	-2
128 (-2)	33k	0.709	61k	294k	0.721	355	126k	0.845	-2	31k	0.695	-2
129 (-3)	31k	0.866	58k	231k	0.676	192	56k	1.033	-3	18k	0.640	-3
152 (-7)	53k	1.490	96k	455k	1.476	364	111k	1.220	-7	39k	0.787	-7
168 (-1)	39k	1.383	98k	263k	0.810	235	86k	1.716	-1	11k	0.632	-1
287 (-3)	38k	1.842	93k	272k	0.828	113	150k	7.902	-3	12k	0.681	-3
329 (-2)	37k	1.737	65k	57k	1.199	65	29k	0.634	-2	6k	0.468	-2
801 (-8)	27k	1.113	55k	202k	0.706	41	61k	1.591	-8	9k	0.405	-8
802 (-3)	16k	1.240	39k	128k	0.950	24	17k	2.195	-3	6k	0.465	-3

In contrast, our method preserves better external features and comparable (if not better) internal features, which leads to better surface reconstructions from the generated medial meshes. It is worth mentioning that, our method does not require large amount of medial spheres, since the design of our framework allows us to directly sample feature spheres without increasing the number of non-feature medial spheres as most of sampling-based and voxel-based methods do. In addition, the Euler characteristic shows the topological correctness and thinness of our generated medial mesh.

6.3 Comparison with the *Voxel Core* Method

We compare the reconstruction quality with the Voxel Core (VC) method [Yan et al. 2018] by setting two voxel sizes: (1) 256^3 with a default pruning parameter $\lambda=0.04$; (2) 512^3 with pruning parameter $\lambda=0.01$ (the default parameter 0.04 prunes excessively so that the medial mesh can not be properly maintained), as shown in Fig. 18. We also show a comparison of two VC results w/o and w/ pruning in Fig. 20.

The medial mesh generated from VC has the following two problems regarding external features. First, the more VC shrinks, the reconstructed shape is more rounded (Fig. 18). If VC shrinks less, redundant structures (Fig. 20 (a)) remain in the resulting medial mesh. Secondly, the boundary curves on VC's medial mesh have good correspondence with the external feature only when the sharp edges of input shape are parallel to the voxels' boundaries. For external features of curves, VC produces zig-zag structures around the feature curves (Fig. 20 (b)). Similar as SAT, VC requires more number of medial spheres (usually \geq 2 times more) than ours in order to generate smooth medial structure around internal features (Fig. 20). Our method, on the contrary, preserves the complete medial structure with a lower reconstruction error and a fewer number of medial spheres (Fig. 20 (c)).

6.4 Ablation Study on Thinning Algorithm

One important property of medial axis is its thinness, *i.e.*, it contains no three-dimensional cells. In this subsection, we give an in-depth

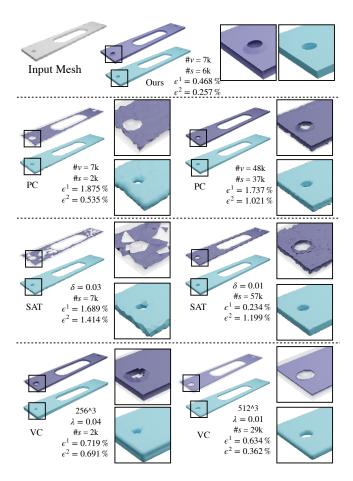


Fig. 18. Qualitative comparison of the medial mesh and quantitative comparison of the reconstructed mesh among our method and two sampling-based methods: PC [Amenta et al. 2001a] and SAT [Miklos et al. 2010], and a voxel-based method VC [Yan et al. 2018].

analysis of the rational of the RPS-based sorting in our geometry-guided thinning algorithm.

The plain thinning algorithm [Ju et al. 2007] prunes tetrahedrons by removing tet-face pairs randomly as long as they are simple pairs (*i.e.*, the face is on the boundary of only one tetrahedron). We found this operation routinely produce open "pockets" (triangles forming open cavities) even though the topology is correct. This is because the choice of simple pairs does not consider the geometry information as some faces are more important than others in a single tetrahedron. An example is shown in Fig. 23. The plain thinning algorithm creates unwanted open "pockets" (Fig. 23 (b)) while our geometry-aware thinning algorithms produces geometrically-accurate medial meshes (Fig. 23 (a)).

The target importance factor σ (in Sec. 4.4) is a parameter that controls the simplicity of the output medial mesh. All tet-faces pairs will be removed no matter what value σ is, and σ plays as a stop sign for removing face-edge pairs. We show the effect of different values of the parameter σ in Fig. 24. For medial mesh with sharp

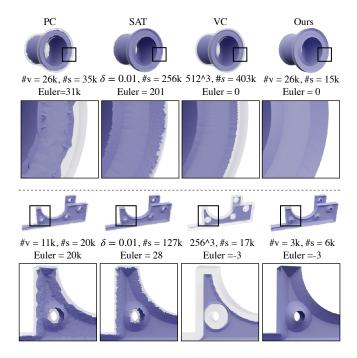


Fig. 19. Comparing the feature preserving ability of our method with two sampling-based methods: PC [Amenta et al. 2001a] and SAT [Miklos et al. 2010], also a voxel-based method VC [Yan et al. 2018].

edges as closed boundaries, a larger value of σ results in a cleaner medial mesh. We use σ = 0.3 for all models in Table. 2. For medial meshes whose boundaries are not closed sharp edges, however, the default value of $\sigma = 0.3$ would remove face-edge pairs that are not suppose to be deleted. To prevent undesired deletion, we use $\sigma = 0.1$ for models whose medial mesh does not ends at sharp features, which includes models in Fig. 26 and models marked ★ in Tab.1 of Supplementary Material.

6.5 More Results

Results on Various Corner Features. We show five example results of our generated medial mesh of various corners in Fig. 22. The first three models (a), (b) and (c) contain saddles incident to both convex and concave edges. The fourth model (d) is a wedge with small angles. The fifth model (e) is a discretized cone with a corner at the tip. Since all of existing methods are known to have difficulties on preserving external corner features, we do not show their results as comparison.

Results on Non-feature Shapes. We show a visual and quantitative comparison of our method with PC [Amenta et al. 2001a], SAT [Miklos et al. 2010] and VC [Yan et al. 2018] on two smooth shapes in Fig. 26. We found that PC, SAT and our method normally give similar reconstruction accuracy while VC has higher errors since it shrank the generated medial mesh during pruning. Similar to models with features, our method generates fewer medial spheres (i.e., PC 20k, SAT 119k, and ours 8k for the bear model) and maintains the thinness property of medial axis while PC and SAT cannot.

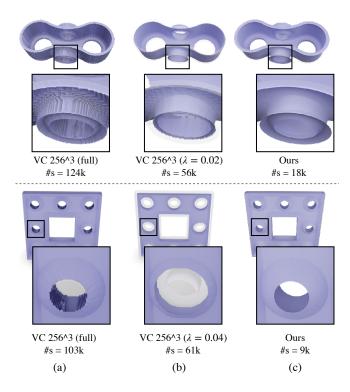


Fig. 20. Comparing medial mesh results of our method (c) with VC [Yan et al. 2018] w/o (a) and w/ (b) λ -pruning.

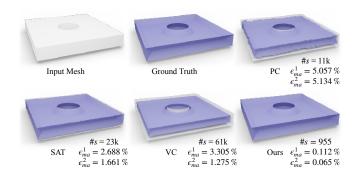


Fig. 21. Quantitative comparison with PC [Amenta et al. 2001a], SAT [Miklos et al. 2010], and VC [Yan et al. 2018] on the medial mesh, with their Hausdorff distances measured w.r.t. the ground-truth medial axis.

We also show a gallery of more results from ABC dataset [Koch et al. 2019] under the 10k/test/2048 folder (see Fig. 25). Please refer to the Supplementary Material for detailed description regarding the selection (Sec. 1) and statistics (Tab. 1).

LIMITATIONS AND FUTURE WORK

It should be noted that our current approach does not guarantee topological preservation for the generated medial mesh w.r.t. the input model, as evidenced from Tab. 1 of Supplementary Material that

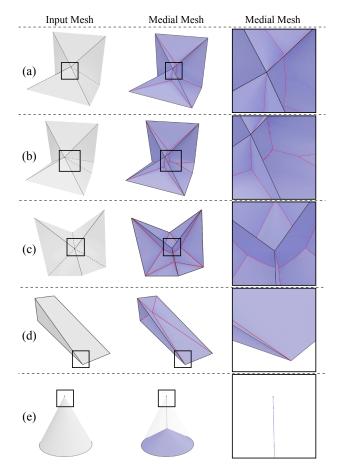


Fig. 22. Examples of corner preservation in our medial mesh: (a) a corner adjacent to four convex edges and two concave edges, forming a saddle; (b) a corner adjacent to four convex edges and three concave edges (two of them are adjacent); (c) a corner adjacent to three convex edges and three concave edges, forming a monkey saddle; (d) a corner adjacent to three sharp edges at small angles, forming a wedge; and (e) a corner at the tip of a discretized cone.

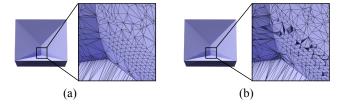


Fig. 23. Comparison of two different thinning strategies on the model of Fig. 22 (a): (a) medial faces sorted by our importance factor α_{ijw} (details in Sec. 4.4 and algorithm provided in Alg. 2 of Supplementary Material; (b) medial faces sorted randomly. We can see that plain thinning strategy (b) creates many open "pockets", while our geometry-guided thinning (a) produces an accurate medial mesh.

there are still 19/73 models having incorrect Euler characteristics. We need to investigate the necessary and sufficient conditions for

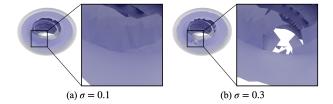


Fig. 24. Thinning results using different target importance factor σ . Default value $\sigma=0.3$ (b) over-pruned the medial mesh which results in holes and detached components, while a smaller value of $\sigma=0.1$ (a) prevents the over-pruning.

topological equivalence under the RPD framework, and come up with some delicate mechanism to preserve the topology of the generated medial mesh. We will leave these topological investigations as our future work.

In addition to the topological preservation issue mentioned above, in this paper we only give experimental evidences but not theoretical proof of correctness for the proposed algorithms, which include: (1) the capability of RPC-based refinement in guaranteeing the topological correctness of internal features; (2) the topological correctness of corner feature preservation; (3) the geometric relationship between different choices of sphere connections (*i.e.*, medial triangles) and their RPS used in our geometry-guided thinning algorithm. More rigorous theoretical guarantees and limitations are needed to be explored in the future.

In addition, the preservation of external features can be effectively evaluated using reconstruction error, however, the same is not obvious for internal features. Moderate errors can be well hidden by the neighboring medial spheres when it comes to shape reconstruction. In our future work, we will consider evaluating the internal features on top of many hex-meshing applications which explicitly rely on internal features to perform solid-meshing of CAD models [Sampl 2000] [Quadros et al. 2004].

8 CONCLUSION

In this paper, we present a novel RPD-based framework for computing the medial axis transform of 3D shapes with preservation of both external and internal medial features. The method is based on the observation that the surface RPC of each medial sphere indicates the set of connected components (CCs) that the sphere has tangential contacts with. Each sphere's CCs can not only be used to update the spheres to their ground truth position and radius, but also tell the information about whether this sphere is on a medial sheet, a seam, or a junction. Such information can be further used to check if the internal or external features are broken in the generated medial mesh, and guide the sphere sampling to preserve those features. Experimental evidences show that our method generates medial meshes with high quality in preserving medial features, both externally and internally. In the future, we believe our RPD-based framework, as a general tool for encoding shapes with features, has potential to be applied to various applications such as shape segmentation [Lin et al. 2022], shape recognition [Hu et al. 2019], and shape deformation [Lan et al. 2020], etc.

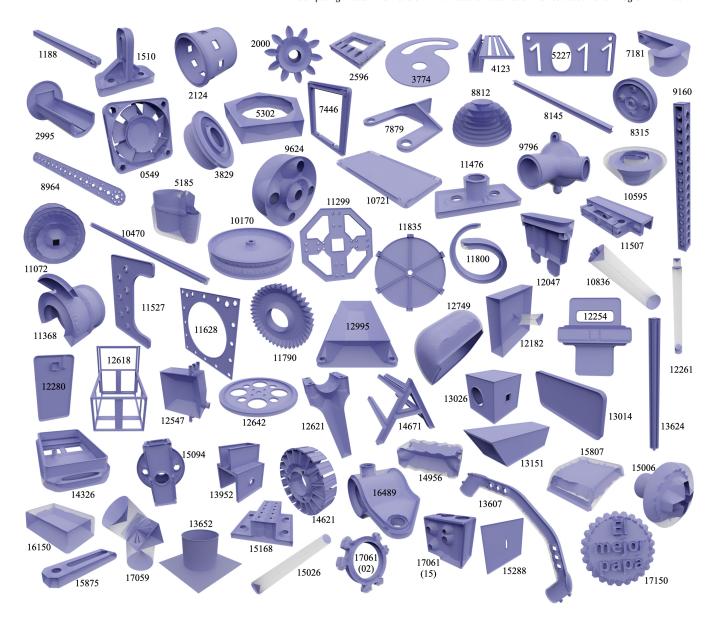


Fig. 25. A gallery of our results using models from ABC dataset under the 10k/test/2048 folder. The selection and statistics details are given in the Supplementary Material in Sec. 1 and Tab. 1.

ACKNOWLEDGMENTS

We would like to thank Shibo Song for helping us prepare figures and video for this paper. Ningna Wang and Xiaohu Guo were partially supported by National Science Foundation (OAC-2007661). Bin Wang was partially supported by National Key Research and Development Program of China (2020YFB1708900).

REFERENCES

Ahmed Abdelkader, Chandrajit L Bajaj, Mohamed S Ebeida, Ahmed H Mahmoud, Scott A Mitchell, John D Owens, and Ahmad A Rushdi. 2020. VoroCrust: Voronoi meshing without clipping. ACM Transactions on Graphics (TOG) 39, 3 (2020), 1-16.

Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. 2001a. The power crust. In Proceedings of the sixth ACM symposium on Solid modeling and applications. 249–266.

Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. 2001b. The power crust, unions of balls, and the medial axis transform. Computational Geometry 19, 2-3 (2001), 127 - 153

Franz Aurenhammer. 1987. Power diagrams: properties, algorithms and applications. SIAM J. Comput. 16, 1 (1987), 78-96.

Gavin Barill, Neil G Dickson, Ryan Schmidt, David IW Levin, and Alec Jacobson. 2018. Fast winding numbers for soups and clouds. ACM Transactions on Graphics (TOG) 37, 4 (2018), 1-12.

Harry Blum et al. 1967. A transformation for extracting new descriptors of shape. Vol. 43. MIT press Cambridge, MA.

Jonathan W Brandt and V Ralph Algazi. 1992. Continuous skeleton computation by Voronoi diagram. CVGIP: Image understanding 55, 3 (1992), 329-338.

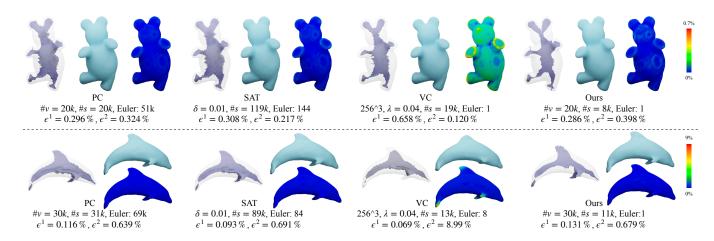


Fig. 26. Comparison with PC [Amenta et al. 2001a], SAT [Miklos et al. 2010], and VC [Yan et al. 2018] on two non-CAD models. We show the generated medial meshes and the reconstructed surfaces, and the color-coded distribution of Hausdorff errors from the reconstructed surface to the input surface.

Frédéric Chazal and André Lieutier. 2005a. Weak feature size and persistent homology: computing homology of solids in \mathbb{R}^n from noisy data samples. In *The twenty-first annual symposium on Computational geometry (SCG '05)*. ACM New York, NY, USA, 255–262.

Frédéric Chazal and André Lieutier. 2005b. The "λ-medial axis". Graphical Models 67, 4 (2005), 304–331.

Yang Chen and Gérard Medioni. 1992. Object Modelling by Registration of Multiple Range Images. *Image and Vision Computing* 10, 3 (1992), 145–155.

Tim Culver, John Keyser, and Dinesh Manocha. 2004. Exact computation of the medial axis of a polyhedron. Computer Aided Geometric Design 21, 1 (2004), 65–98.

Tamal K Dey, Hyuckje Woo, and Wulue Zhao. 2003. Approximate medial axis for CAD models. In Proceedings of the eighth ACM symposium on Solid modeling and applications. 280–285.

Tamal K Dey and Wulue Zhao. 2002. Approximate medial axis as a voronoi subcomplex. In Proceedings of the seventh ACM symposium on Solid modeling and applications. 356–366

Tamal K Dey and Wulue Zhao. 2004. Approximating the medial axis from the Voronoi diagram with a convergence guarantee. Algorithmica 38, 1 (2004), 179–200.

Zhiyang Dou, Cheng Lin, Rui Xu, Lei Yang, Shiqing Xin, Taku Komura, and Wenping Wang. 2022. Coverage Axis: Inner Point Selection for 3D Shape Skeletonization. In Computer Graphics Forum, Vol. 41. Wiley Online Library, 419–432.

Noura Faraj, Jean-Marc Thiery, and Tamy Boubekeur. 2013. Progressive medial axis filtration. In SIGGRAPH Asia 2013 Technical Briefs. 1–4.

Peter Giblin and Benjamin B Kimia. 2004. A formal classification of 3D medial axis points and their local geometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26, 2 (2004), 238–251.

Jianwei Hu, Bin Wang, Lihui Qian, Yiling Pan, Xiaohu Guo, Lingjie Liu, and Wenping Wang. 2019. MAT-Net: Medial Axis Transform Network for 3D Object Recognition. In Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI'19). 774–781.

Tao Ju, Qian-Yi Zhou, and Shi-Min Hu. 2007. Editing the topology of 3D models by sketching. ACM Transactions on Graphics (TOG) 26, 3 (2007), 42-es.

Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. 2019. Abc: A big cad model dataset for geometric deep learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 9601–9611.

Lei Lan, Ran Luo, Marco Fratarcangeli, Weiwei Xu, Huamin Wang, Xiaohu Guo, Junfeng Yao, and Yin Yang. 2020. Medial Elastics: Efficient and Collision-Ready Deformation via Medial Axis Transform. ACM Trans. Graph. 39, 3, Article 20 (apr 2020).

Pan Li, Bin Wang, Feng Sun, Xiaohu Guo, Caiming Zhang, and Wenping Wang. 2015.
Q-mat: Computing medial axis transform by quadratic error minimization. ACM Transactions on Graphics (TOG) 35, 1 (2015), 1–16.

Cheng Lin, Lingjie Liu, Changjian Li, Leif Kobbelt, Bin Wang, Shiqing Xin, and Wenping Wang. 2022. SEG-MAT: 3D Shape Segmentation Using Medial Axis Transform. IEEE transactions on visualization and computer graphics 28, 6 (2022), 2430–2444.

Lu Liu, Erin W Chambers, David Letscher, and Tao Ju. 2010. A simple and robust thinning algorithm on cell complexes. In *Computer Graphics Forum*, Vol. 29. Wiley Online Library, 2253–2260. Jaehwan Ma, Sang Won Bae, and Sunghee Choi. 2012. 3D medial axis point approximation using nearest neighbors and the normal field. The Visual Computer 28, 1 (2012), 7–19

Balint Miklos, Joachim Giesen, and Mark Pauly. 2010. Discrete scale axis representations for 3D geometry. In ACM SIGGRAPH 2010 papers. 1–10.

Victor Milenkovic. 1993. Robust Construction of the Voronoi Diagram of a Polyhedron... In CCCG, Vol. 93. Citeseer, 473–478.

Yiling Pan, Bin Wang, Xiaohu Guo, Hua Zeng, Yuexin Ma, and Wenping Wang. 2019.
Q-mat+: An error-controllable and feature-sensitive simplification algorithm for medial axis transform. Computer Aided Geometric Design 71 (2019), 16–29.

Stephen M Pizer, Kaleem Siddiqi, Gabor Székely, James N Damon, and Steven W Zucker. 2003. Multiscale medial loci and their properties. *International Journal of Computer Vision* 55, 2 (2003), 155–179.

WR Quadros, K Ramaswami, FB Prinz, and B Gurumoorthy. 2004. LayTracks: a new approach to automated geometry adaptive quadrilateral mesh generation using medial axis transform. International journal for numerical methods in engineering 61, 2 (2004), 209–237.

Punam K Saha, Gunilla Borgefors, and Gabriella Sanniti di Baja. 2016. A survey on skeletonization algorithms and their applications. *Pattern recognition letters* 76 (2016), 3–12.

Peter Sampl. 2000. Semi-structured mesh generation based on medial axis. In 9th International Meshing Roundtable. Citeseer.

Evan C Sherbrooke, Nicholas M Patrikalakis, and Erik Brisson. 1996. An algorithm for the medial axis transform of 3D polyhedral solids. *IEEE transactions on visualization* and computer graphics 2, 1 (1996), 44–61.

Kaleem Siddiqi and Stephen Pizer. 2008. Medial representations: mathematics, algorithms and applications. Vol. 37. Springer Science & Business Media.

André Sobiecki, Andrei Jalba, and Alexandru Telea. 2014. Comparison of curve and surface skeletonization methods for voxel shapes. Pattern Recognition Letters 47 (2014), 147–156.

Andrea Tagliasacchi, Thomas Delame, Michela Spagnuolo, Nina Amenta, and Alexandru Telea. 2016. 3d skeletons: A state-of-the-art report. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 573–597.

Dong-Ming Yan, Bruno Lévy, Yang Liu, Feng Sun, and Wenping Wang. 2009. Isotropic remeshing with fast and exact computation of restricted Voronoi diagram. In Computer graphics forum, Vol. 28. Wiley Online Library, 1445–1454.

Yajie Yan, David Letscher, and Tao Ju. 2018. Voxel cores: Efficient, robust, and provably good approximation of 3d medial axes. ACM Transactions on Graphics (TOG) 37, 4 (2018), 1–13.

Yajie Yan, Kyle Sykes, Erin Chambers, David Letscher, and Tao Ju. 2016. Erosion thickness on medial axes of 3D shapes. ACM Transactions on Graphics (TOG) 35, 4 (2016), 1–12.