

Coverage Axis: Inner Point Selection for 3D Shape Skeletonization

Zhiyang Dou¹ , Cheng Lin² , Rui Xu³ , Lei Yang¹ , Shiqing Xin³ , Taku Komura¹ , Wenping Wang⁴ 

¹ The University of Hong Kong

² Digital Content Technology Center, CROS, Tencent Games

³ Shandong University

⁴ Texas A&M University




Figure 1: A gallery of 3D shape skeletonization results by Coverage Axis from surface mesh input.

Abstract

In this paper, we present a simple yet effective formulation called Coverage Axis for 3D shape skeletonization. Inspired by the set cover problem, our key idea is to cover all the surface points using as few inside medial balls as possible. This formulation inherently induces a compact and expressive approximation of the Medial Axis Transform (MAT) of a given shape. Different from previous methods that rely on local approximation error, our method allows a global consideration of the overall shape structure, leading to an efficient high-level abstraction and superior robustness to noise. Another appealing aspect of our method is its capability to handle more generalized input such as point clouds and poor-quality meshes. Extensive comparisons and evaluations demonstrate the remarkable effectiveness of our method for generating compact and expressive skeletal representation to approximate the MAT.

CCS Concepts

• **Computing methodologies** → **Shape analysis**;

1. Introduction

With the capability of effectively capturing the underlying structures of 3D shapes, skeletal representations have been an important tool in various applications of shape analysis and geometric processing, such as 3D reconstruction [WHG*15; THP*19; ACK01], volume approximation [SKS11; SCYW13], shape segmentation [LLL*20], shape abstraction [DXX*20], pose estimation [SFC*11; YRL*21] and animation [BP07; YYG18], etc.

Curve skeletons, consisting of only 1D curves, have been extensively researched [ATC*08; MWO03; TAOZ12; XZKS19] due to

their simplicity and intuitiveness. Curve skeletons are usually empirically understood, while Dey and Sun [DS06] give a mathematical definition based on Medial Geodesic Function. Generally, the curve representation only applies to tubular components instead of arbitrary shapes, which thus cannot be considered as a generalized tool for shape analysis.

Another prominent example of skeletal representations is called the Medial Axis Transform (MAT) [Blu*67]. The MAT is defined by a union of the maximally inscribed balls inside the shape with the associated radius functions. Different from the curve skeleton,

the MAT has a consistent definition for arbitrary shapes. In addition to curves, the MAT consists of both curve-like and surface-like structures thus leading to significantly better representational ability. That being said, the MAT is difficult to use which mainly manifests in two aspects. First, the MAT is notoriously sensitive to boundary noise, i.e., small perturbations on the boundary surface will result in dramatic changes on the medial axis. Second, the computation of the MAT usually relies on stringent requirements of the input geometry, such as the watertightness and manifoldness of the surface.

In this paper, we propose Coverage Axis, a novel and simple formulation to generate skeletal representations for 3D shapes. Our goal is to give a compact approximation of the MAT, while this approximation should inherit good geometric and topological properties of the MAT but overcome its aforementioned drawbacks.

We observe a medial sphere can be considered as an abstraction of local geometry, while the union of all the local geometries forms the entire shape. With this insight, our key idea is to formulate a Set Cover Problem, of which goal is to identify the smallest sub-collection of whose union equals the universe. Specifically, we aim to find the minimum number of dilated inner balls that cover sub-regions on the surface (i.e., sub-collection) to approximate the whole shape (i.e., universe).

This formulation inherently induces a compact and expressive representation that approximates the MAT. First, the coverage constraint enforces consistency with the shape structure. Second, selecting the fewest possible candidates for shape approximation not only gives a simplified representation, but also favors the interior points that dominate larger area, which corresponds to the definition of MAT, i.e., maximally inscribed spheres.

Compared to the existing methods that mainly focus on MAT simplification [ACK01; LWS*15; GMPW09; FLM03; DZ02; SFM07; MGP10], our method has several appealing aspects. First, its formulation is simple yet effective, which does not require any complex geometric processing nor a computationally-costly pipeline to derive skeletal points. Compared to metrics like QEM [GH97] and Spherical QEM [TGB13] which are used to construct the medial axes by only considering the local geometry, our set coverage formulation jointly leverages all the surface points and inner points for computing MAT approximation, which results in considering the global shape structure, and thus leads to a better abstraction of the overall shape as well as a shape-aware point distribution. More importantly, we can handle the input of poor-quality meshes or point clouds since our method is based on the coverage of point sets, while other algorithms usually rely on watertight surfaces with decent quality. Finally, our method takes as input a set of overfilled inner point candidates and selects the most expressive ones as the skeletal points; thus it does not necessarily need to compute a rigorously defined MAT, but can work with randomly generated points inside the input.

We demonstrate the effectiveness and robustness of our method on a variety of 3D shapes. The extensive evaluation results reveal that our novel and simple formulation for inner point selection effectively captures the global structure and the fundamental geometry of input shapes, thus providing informative skeletal representations to approximate the structure of the MAT.

2. Related Work

Curve skeletonization Curve skeletonization has been extensively researched in computer vision and computer graphics. Traditional methods rely on hand-crafted rules to utilize geometric information for curved skeleton computation [AM96; MWO03; SLSK07; ATC*08; NBPF11; TAOZ12; XZKS19; CZC*20]. Specifically, Ma et al. [MWO03] applies using radial basis functions (RBFs) for skeleton extraction. Sharf et al. [SLSK07] adopt a deformable model evolution that captures the object's volumetric shape and then generates the approximation for the curve skeleton. Au et al. [ATC*08] compute the skeleton via mesh extraction. Livesu et al. [LGS12] reconstruct the curve skeletons of 3D shapes using the visual hull. Mean Curvature Skeleton [TAOZ12] formulates the skeletonization problem via mean curvature flow, which drives the curvature flow towards the extreme so as to collapse the input mesh. The mesh collapse also induces an intermediate result called meso-skeleton that consists of both surface-like and curve-like structures. Recently, Cheng et al. [CZC*20] propose a method for skeletonization using the dual of shape segmentation.

At the same time, learning-based methods [XZKS19] have been proposed for predicting curve skeletons. Although curve skeletons are able to represent tubular shapes, it has issues representing arbitrary shapes, e.g., shapes with flat components.

Medial axis transform As a more general skeletal representation, medial axis transform (MAT) [Blu*67] is able to encode arbitrary shapes with curve-like and mesh-like structures. The maximal balls defined in the volume together with their locii complete the representation of the shape.

The most commonly used technique for extracting MAT is to initialize the medial surface using the Voronoi diagrams. Then a simplified medial surface is achieved by applying optimization on spikes pruning or mesh tessellations based on various rules. Following this way, there are many existing approaches. Angle-based filtering methods [ACK01; FLM03; DZ02; SFM07] reach a simplification based on the angle formed by the point of MA and its two closest points on the boundary. These methods usually produce a simplified result with local features well preserved; however, they suffer from preserving the original topology of the objects. λ -medial axis methods [CL05; CCT11] are another series of methods to simplify the computation of MAT, which adopt cumradius of the closest points of a medial point as a pruning criterion. The main drawback of these methods is poor feature preserving ability at different scales [ABE09].

Compared with the aforementioned methods, Scale Axis Transform (SAT) [MGP10] prunes spikes more effectively. Unstable medial axis points are identified if the corresponding medial balls are covered by the neighboring medial balls during the multiplicative growing. However, SAT typically has a significant computational cost and tends to destroy the topology by introducing new topological structures at large scales. To overcome these issues, progressive medial axis filtration (PMAF) [FTB13] proposes to perform successive edge collapse based on sphere absorption to preserve the topology as well as improve time efficiency. Our idea also uses a dilated medial balls; we pay special attention to the relationship between inner balls and surface samples. Different from SAT that

favors a dense representation with a large number of vertices on the medial axis, our method is able to yield a compact skeletal representation without destroying the shape structure.

Meanwhile, other strategies have been proposed to achieve simplified MAT, e.g., Delta Medial Axis (DMA) [MLM16], Bending Potential Ratio (BPR) pruning [SBH*11], Erosion Thickness (ET) measure [YSC*16], voxelization based λ -pruning [YLJ18]. The simplified MAT, without doubt, cannot reconstruct the original shape exactly, and thus Sun et al. [SCYW13] propose to control the approximation error by computing the Hausdorff distance. Rebain et al. [RLS*21] introduce medial fields derived from the MAT to represent the original shape using learning techniques.

To achieve high accurate approximation of the original shape, Li et al. [LWS*15] propose Q-MAT that adopts quadratic error metric (QEM) [GH97] for MAT simplification. As an edge collapse method, Q-MAT is fast and produces a piecewise linear approximation of the MAT. Q-MAT collapses edges based solely on local information, which sometimes leads to unreasonable spatial distribution of vertices. Also, Q-MAT is limited to 2-manifold surface scope. Additionally, Q-MAT heavily relies on a good MAT initialization; when the initial MA is low quality, Q-MAT can not produce correct results by edge collapse. Instead, we tackle this problem from a global view, which can preserve representative inner points with relatively even spatial distribution. As our method is based on the coverage of point sets, we can handle inputs of poor-quality meshes or point clouds. Furthermore, our method, as we will demonstrate, has good adaptability to low-quality candidate points. We refer readers to [TDS*16] for a detailed survey covering various forms of skeletons.


Point cloud skeletonization Nowadays, skeletonization of point cloud is drawing people’s attention due to the easy availability of point cloud data [TZC09; CTO*10; LGS12; HWC*13; RAV*19; WHG*15; YYW*20]. In particular, L1-medial skeleton [HWC*13] contracts point clouds based on locally optimal projections (LOP) [LCLT07]. LSMAT [RAV*19] takes a densely sampled oriented point set as input and computes an MAT approximation on the basis of Signed Distance Function (SDF). Wu et al. [WHG*15] associates the surface points to the inner point residing on the meso-skeleton [TAOZ12]. However, represented by unstructured points, their skeletonization lacks topological constraints. Additionally, the reconstruction results of [WHG*15] are still with large error.

With the success of deep neural networks, learning-based methods have been proposed for predicting skeletons of point clouds [YYW*20; LLL*21]. P2MAT-NET [YYW*20] transforms sparse point cloud to an output set of spheres to approximate the medial balls. Point2Skeleton (P2S) [LLL*21] learns skeletal representations from generalized point clouds in an unsupervised manner. Learning-based methods do not guarantee accurate computation of geometric features. Furthermore, they often suffer from generalization ability due to the dependency on the training data.

3. Preliminaries

3.1. Medial Surface

Given a closed, oriented and bounded two-manifold surface S in \mathbb{R}^3 , medial axis is defined as the locus of the center of maximally inscribed spheres that are tangent to S at two or more points. The medial axis \mathcal{M} of S , together with its radius function \mathcal{R} , forms the medial axis transform (MAT), denoted by a pair $(\mathcal{M}, \mathcal{R})$.



3.2. Voronoi Diagram & Power Diagram




Figure 2: 2D Voronoi diagram and Power diagram (solid lines) and their dual graphs: Delaunay triangulation and Regular triangulation (dashed lines). The weight of each seed in the Power diagram is visualized by a disk.

Voronoi diagram is a partition of the domain $\Omega \subset \mathbb{R}^d$ into regions, close to a set of points called generators $\{\mathbf{x}_i \in \Omega\}_{i=1}^n$. Each region, also called cell, is defined as

$$\Omega_i^{\text{vor}} : \{\mathbf{x} \in \Omega \mid \|\mathbf{x} - \mathbf{x}_i\| \leq \|\mathbf{x} - \mathbf{x}_j\|, j \neq i\}.$$

A well-known MAT initialization technique is to use the Voronoi diagram inside a model generated by sampled points on the surface.

Power diagrams [Aur87] can be viewed as an extension of Voronoi diagrams, where each generator \mathbf{x}_i is equipped with a weight w_i to control its influence. By defining the power distance $d^{\text{pow}}(x, x_i)$ between x and the weighted generator x_i to be $\|\mathbf{x} - \mathbf{x}_i\|^2 - w_i$, the cell associated with \mathbf{x}_i is defined by

$$\Omega_i^{\text{pow}} : \{\mathbf{x} \in \Omega \mid d^{\text{pow}}(x, x_i) \leq d^{\text{pow}}(x, x_j), j \neq i\}.$$

A generator with a larger weight is more dominant, and when all the weights are equal, the Power diagram is reduced to a Voronoi diagram. In [ACK01], the approximate MAT structure is built based on regular triangulation, the dual of Power diagram, with the weights as squared radius of inside and outside poles (a subset of Voronoi vertices). Examples of a Voronoi diagram and a Power diagram are shown in Figure 2 respectively.

4. Method

Our goal is to generate a compact skeletal representation of a given 3D shape that approximates its MAT, for capturing the fundamental structure of the shape.

Our method consists of three main steps. We first generate a set of inner points as the candidate skeletal points. Then, we introduce a novel formulation called coverage axis, which is based on the set cover problem, to find a minimal sub-collection of candidates to recover the input shape. Finally, we analyze the connectivity of the selected points to form a connected structure with edges and triangles.




Figure 3: Pipeline of inner point selection. (a) Input 3D model. (b) Inner point candidates. (c) Dilated inner balls. (d) Selected dilated balls based on surface point coverage. (e) Selected inner balls with original radius. (f) Selected inner points.

4.1. Inner Points Generation

Our method starts with the generation of candidate points located inside a given 3D shape. The input to our system can be either surface meshes or point clouds with normals.

For a mesh input, following a common initialization technique, we compute the Voronoi diagram of a set of points sampled on the mesh and extract the inner points, which become the candidate points of the skeleton. The inside Voronoi diagram becomes the initial coverage axis structure.

For a point cloud input (with normals), similar to a mesh input, we first compute the Voronoi diagram w.r.t. surface samples. Then, we extract the inner points as follows. We first compute its Delaunay triangulation which is the dual of the Voronoi diagram w.r.t. the input point cloud. Consider a Voronoi vertex p and the vertices of its dual tetrahedron p'_0, p'_1, p'_2 and p'_3 . The candidate p is considered as an inner point only if we have $\vec{pp}'_i \cdot \vec{n}(p'_i) > 0, \forall i = 0, 1, 2, 3$. Here $\vec{n}(p'_i)$ is the input normal of p'_i and \cdot is dot product. See more details in Appendix A. Note other inside-outside query methods such as [BDS*18] are also applicable.

The Voronoi-based point generation is based on principled geometric transform which can generate candidates with good quality. Nevertheless, note our method does not rely on strictly defined MAT for initialization. We can also randomly sample a set of points inside the shape and use them as skeleton candidates (see Sec. 6.1 for a detailed discussion). This property significantly enhances the flexibility of our method for handling various inputs in different applications.

4.2. Point Selection Based on Set Coverage

Given a set of elements (called universe) and its subsets whose union equals the universe, the set cover problem (SCP) is to identify the minimum sub-collection of these subsets whose union equals the universe. Inspired by this formulation, we now introduce our Coverage Axis. In the last step, we obtain a set of inner point candidates $P = \{p_i\}$. We estimate the radii of these points $R = \{r_i\}$ as their closest distances to the boundary surface, resulting in candidate balls. Now we slightly dilate all the balls by adding a small value to their radii, i.e., $r'_i = r_i + \delta_r$, leading to a set of dilated balls $B = \{(p_i, r'_i)\}$. The dilation makes our algorithm robust to insignificant details and noise, for which we give detailed discussions later in Sec. 6.2.

Our goal is to find the minimum number of dilated balls that cover all the sampled points $S = \{s_j\}$ on the surface. For this purpose, we introduce a coverage matrix $\mathbf{D} \in \{0, 1\}^{m \times n}$, where m and

n are total numbers of sampled surface points and candidate skeletal points, respectively. Each element $d_{ji} \in \{0, 1\}$ of \mathbf{D} indicates if a surface point s_j is covered by the dilated ball (p_i, r'_i) :

$$d_{ji} = \begin{cases} 1, & \text{if } \|p_i - s_j\|_2 \leq r'_i, j = 1, \dots, m, i = 1, \dots, n \\ 0, & \text{if } \|p_i - s_j\|_2 > r'_i, j = 1, \dots, m, i = 1, \dots, n. \end{cases} \quad (1)$$

Let $\mathbf{v} \in \{0, 1\}^{n \times 1}$ be a decision vector, where the i -th element $v_i \in \{0, 1\}$ indicates if a candidate skeletal point p_i is selected. Now we can derive the formulation of our 0-1 integral optimization problem:

$$\begin{aligned} \min & \|\mathbf{v}\|_2 \\ \text{s.t.} & \mathbf{D}\mathbf{v} \geq \mathbf{1}, \end{aligned} \quad (2)$$

where $\mathbf{1}$ is a vector of ones and \geq is applied element-wise to the vector entries. Here we minimize the norm of the decision vector \mathbf{v} under a certain constraint. The constraint enforces each sampled point to be covered by the selected skeletal balls, which ensures consistency between the selected skeleton and the input shape. With this simple formulation, our method allows a global consideration of the overall shape structure and effectively selects the most predominant inner candidates, giving an expressive and compact skeletal representation to faithfully capture the fundamental geometry and topology. We solve Eq. 2 using Mixed-integer linear programming (MILP) solver in Matlab (MathWorks 2021). The pipeline of inner point selection is demonstrated in Figure 3. We further discuss a series of properties of our method, e.g., centrality, robustness to noise, etc., in Sec. 6.

It is notable that the constraint in Eq. 2 does not always lead to a solution. For example, if there are points not covered by any dilated balls (assume the dilation is tiny), no solution can be found. Nevertheless, this barely happens and the fixed parameters always lead to a solution for the extensive experiments as we will show later. This is owing to the proper initialization for generating inside candidates: Voronoi Diagram or densely sampled points. Voronoi Diagram produces balls maximally inscribed to the surface, thus already giving tight-fitting candidates. The randomly generated inner points are considerably dense, of which a subset approximates the vertices of the Voronoi Diagram, and the excessive balls provide abundant candidates for solving SCP.

4.3. Connection Establishment

Once the skeletal points are determined, in this step, we analyze their connectivity to form a structured mesh. Generally, if the input is a surface mesh, the topological connections of the skeletal points




Figure 4: Shape approximation results of Coverage Axis. First row: the input shapes. Second row: the output skeletal representations for MAT approximation. Third row: the reconstructed shapes by the skeletal representations. Fourth row: color coding of the reconstruction error between each point in the reconstruction and the input shape, defined relative to the diagonal length of the bounding box.

are easier to be derived from the surface. However, this is difficult when the input is a point cloud that only consists of unorganized points.

Considering our method flexibly allows different input types (i.e., surface mesh and point cloud) with different candidate generation strategies (i.e., Voronoi-based and random sampling), we propose to use different ways to build the connections of skeletal points. More details can be found in Appendix B.

Mesh with candidates from Voronoi diagram

This configuration usually gives the best skeletonization quality since not only the boundary surface is structured, but the inner points are derived from principled geometric transform, thus also having a good structure. Accordingly, we can obtain selected points embedded on the Voronoi diagram (the part locating inside the model). Note that such a Voronoi diagram provides a correct topological connection w.r.t. the input 3D model. We then can simplify this structure using off-the-shelf method with selected points. In this paper, we follow the edge collapse strategy of Q-MAT [LWS*15]. But the key difference is that the result selected by set coverage serves as a series of anchors; these points, like nails, are guiding the direction of further simplification so as to ensure a more reasonable distribution w.r.t. original geometry. In particular, for an edge to be collapsed, if one endpoint is not the selected point but another is, we merge the point to the selected point; if two endpoints are both selected points, we preserve this edge and



skip to the next edge; if both two endpoints are not selected, we merge them to the optimal contraction target same as [LWS*15]. Meanwhile, topology preservation [DEGN98] and mesh inversion avoidance algorithm [GH97] are applied during the collapse.

More generalized configurations This case takes the input of either mesh or point cloud, while the candidate skeletal points can be generated by Voronoi diagrams or random sampling. Different from the first case, where both inner and outside structures are well defined, it is challenging to derive precise topological connections of the skeletal points. We seek for an approximation of the MAT connection similar to [ACK01] as follows.

Suppose P^+ is the set of all selected points by set coverage. Let S denotes the set of the surface. We then compute the Power diagram **PD** on the set of $\{B^l(p_i, r'_i), p_i \in P^+\} \cup \{B(s_j, \delta_r), s_j \in S\}$ and establish the connection by extracting edges among the selected inner points on the dual of **PD**, regular triangulation **RT**. In order to suppress over-connection, we up-sample on the input point cloud following [HWG*13], leading to 30000 surface samples. These points are treated as surface points in **RT** during connection establishment. Note that r' can be adjusted accordingly during the computation of **RT** to achieve flexible control over the connection.

5. Experimental Results

In this section, we conduct qualitative and quantitative evaluations of the proposed method to comprehensively demonstrate its effectiveness. We implement and experiment on a computer with a 4.00 GHz Intel(R) Core(TM) i7-6700K CPU and 16 GB memory. The size of all the models are normalized to the [0, 1] range.

All the experimental results are generated using a consistent parameter setting, i.e., for all kinds of input, we always set $\delta_r = 0.02$. For a mesh input, we sample $|C| = 4000$ points from the surfaces. The resulting set is denoted by C and is used to compute the Voronoi diagram, whose vertices inside the input mesh serve as candidate inner points set, denoted by P . Then we sample $|S| = 1500$ surface points as the set S to be covered in the SCP. Here $|S|$ is set to be less than $|C|$ to reduce the computational cost (see Sec. 6.5 for a detailed evaluation). The input point clouds all have 2000 points, and thus we have $|S| = |C| = 2000$. Detailed analyses on these parameters δ_r, S, C and P are given in Sec. 6.4.

As for evaluation metrics, besides one-sided Hausdorff distance (HD) from the input surface to the surface reconstructed by the skeleton that is employed by Q-MAT [LWS*15], we adopt two-sided Hausdorff distance since the former is not informative to fully reveal the consistency between two surfaces. We use $\vec{\epsilon}$ to denote the HD from surface to reconstruction, $\overleftarrow{\epsilon}$ to denote the HD from reconstruction to surface, and $\overleftrightarrow{\epsilon}$ to represent two-sided HD.

We show a set of qualitative results on various 3D shapes in Figure 1. The shape approximation results are shown in Figure 4. The reconstruction results are obtained by interpolating the medial balls based on the skeleton connectivity similar to [LWS*15]. Please refer to the Appendix for more implementation details.

5.1. Coverage Axis from Mesh Input

We conduct comparisons with representative MAT simplification methods including Q-MAT [LWS*15] and SAT [MGP10]. The errors are scaled by the length of the diagonal of the input model bounding box.

Comparison with Q-MAT [LWS*15] For a fair comparison, the number of vertices in the medial axis are set the same for both methods. The quantitative results are reported in Table 1. It can be seen that our method is comparable with or better than Q-MAT in terms of reconstruction error. Additionally, we conduct comparisons with Q-MAT on the robustness against the input noise, whose results are given in Sec. 6.2.

From the experiments, it is revealed that the problem with Q-MAT is that it may yield an uneven distribution of inner points in parts where the geometry is similar (Figure 5 (a) first and second row). This is because the edge collapse strategy of Q-MAT is solely based on local approximation errors, which lacks the global constraints of the whole geometry. This problem is even worse once a highly simplified representation is needed.

In contrast, our method effectively leverages the global geometric and topological features of the input shape for skeleton generation. The point selection step jointly considers the overall input points and candidates, while the selected expressive points are used as fixed anchors to preserve the global topology. Hence, the skeletal representations generated by our method are equipped with shape-aware point distributions as well as connectivity.

It is worth noting that Q-MAT heavily relies on the quality of the input mesh and the initial Voronoi diagram to compute the medial axis skeleton. In contrast, our method does not have this restriction.

Table 1: Quantitative comparison on shape approximation error between Q-MAT and Coverage Axis.

Model	V	Q-MAT			Coverage Axis		
		$\vec{\epsilon}$	$\overleftarrow{\epsilon}$	$\overleftrightarrow{\epsilon}$	$\vec{\epsilon}$	$\overleftarrow{\epsilon}$	$\overleftrightarrow{\epsilon}$
Ant-1	46	2.011%	2.236%	2.236%	2.894%	2.601%	2.894%
Armodillo	95	4.178%	2.125%	4.178%	2.947%	2.793%	2.947%
Bird	54	0.954%	1.539%	1.539%	1.682%	1.816%	1.816%
Bunny	106	2.127%	2.596%	2.596%	2.526%	2.496%	2.526%
Chair-1	79	1.413%	1.534%	1.534%	2.246%	1.976%	2.246%
Crab	58	1.422%	1.794%	1.794%	1.458%	1.759%	1.759%
Cup	132	1.213%	2.481%	2.481%	2.599%	2.636%	2.636%
Desk	151	1.068%	1.930%	1.930%	2.273%	2.164%	2.273%
Dog	49	2.582%	3.266%	3.266%	2.077%	1.850%	2.077%
Fertility	79	3.016%	4.259%	4.259%	2.703%	2.377%	2.703%
Fish	53	2.336%	1.764%	2.336%	2.325%	2.281%	2.325%
Hand-1	42	1.594%	1.780%	1.780%	1.334%	1.750%	1.750%
Hand-2	60	1.203%	1.747%	1.747%	1.715%	1.716%	1.716%
Human-1	45	3.231%	1.929%	3.231%	2.729%	2.644%	2.729%
Human-2	46	1.614%	1.572%	1.614%	1.498%	1.526%	1.526%
Human-3	61	2.265%	1.561%	2.265%	1.788%	1.707%	1.788%
Kitten	45	2.490%	2.309%	2.490%	2.875%	2.893%	2.893%
Octopus-1	49	1.811%	1.845%	1.845%	1.825%	1.836%	1.836%
Octopus-2	68	2.356%	1.765%	2.356%	2.256%	2.316%	2.316%
Pig	56	2.503%	2.420%	2.503%	2.107%	2.183%	2.183%
Plane	53	1.108%	1.184%	1.184%	1.430%	1.424%	1.430%
Snake	39	1.183%	1.639%	1.639%	1.335%	1.623%	1.623%
Spectacle	45	2.191%	1.061%	2.191%	2.133%	1.542%	2.133%
Bear	32	2.348%	2.619%	2.619%	3.185%	3.229%	3.229%
Vase	106	2.906%	2.453%	2.906%	3.183%	2.954%	3.183%
Average	-	2.045%	2.056%	2.341%	2.205%	2.164%	2.264%

$|V|$ The number of skeletal points.

$\vec{\epsilon}$ One-sided HD from surface to reconstruction.

$\overleftarrow{\epsilon}$ One-sided HD from reconstruction to surface.

$\overleftrightarrow{\epsilon}$ Two-sided HD between original surface and reconstruction.

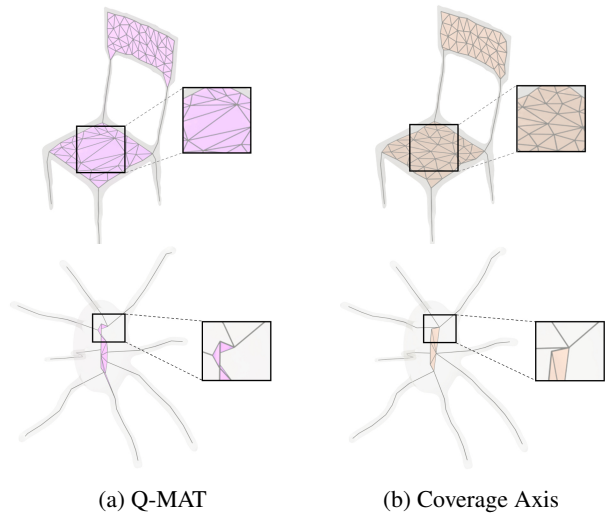


Figure 5: Comparison between Q-MAT and Coverage Axis on the quality of the generated medial surfaces. Note that the number of vertices is same.




Figure 6: Comparison with Q-MAT for poor-quality mesh that can fail the construction of Voronoi Diagram.

Figure 6 shows a group of examples where Q-MAT yields incorrect structures due to the failure of initialization caused by the poor mesh quality, where our algorithm successfully gives the proper skeletal representation.

Comparison with SAT [MGPI0] One of the main setbacks of this method is that SAT favors a dense representation with a large number of vertices, thus incapable of generating a simple and compact skeleton for shape abstraction. In this comparison with SAT, we test different values of its scaling factor, i.e., $\alpha = 1.1, 1.5$ and 2.0 . The qualitative and quantitative comparison results are shown in Figure 8. On the one hand, despite high accuracy yielded by SAT using smaller α , its representation has a large number of vertices (Figure 8 (a-c)). This redundant representation is not amenable for various applications that require skeleton to be simple and expressive, e.g., shape matching and retrieval [SSGD03], skeletal animation [MLT88], animated mesh approximation [YYG18; TGBE16], etc. On the other hand, using a larger α for a higher abstraction level destroys the shape structure and causes large approximation errors. Our method, on the contrary, provides trade-offs across fidelity, efficiency and compression abilities.

Both Q-MAT and SAT can only excel at handling watertight surface meshes. As aforementioned, the surface mesh input allows robust computation of the Voronoi diagram, which offers good structural information from both inside and outside. Instead, generating skeletal representations for point clouds is more challenging and beyond the capacity of these methods. Our method is able to take unstructured point sets as input and generate skeletal representations of which connected structures are recovered. We will evaluate this property of our method in the following section.

5.2. Coverage Axis from Point Cloud Input

For point cloud input, we compare with two closely relevant methods: Deep Point Consolidation (DPC) [WHG*15] and Point2Skeleton [LLL*21], given that these two methods can also predict MAT-based skeletons directly from point clouds. Note Point2Skeleton is a deep learning-based method that requires training from a large amount of data.

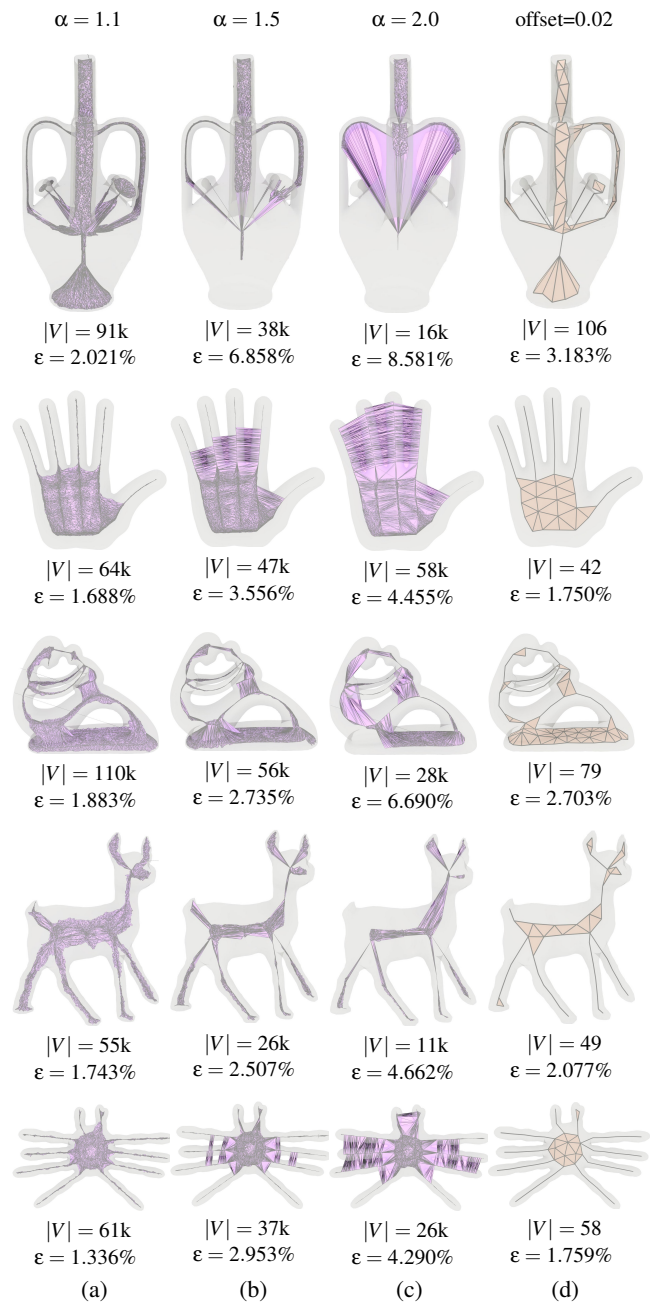


Figure 8: Qualitative comparison with SAT. (a-c) SAT with mesh input. (d) Coverage Axis with mesh input. Note that Coverage Axis (ours) achieves trade-offs across fidelity, efficiency and compression abilities compared with SAT.

Comparison with Deep Point Consolidation [WHG*15] As shown in Figure 7 and Table 2, this method can only produce unstructured inner points without connections. Although the skeletons generated by DPC have denser points, the reconstruction results still exhibit large errors. Even worse, the reconstructed topologies by deep point consolidation are usually inconsistent with the original input. In contrast, our reconstruction results reach better ap-




Figure 7: Comparison with existing shape skeletonization methods for point clouds.

proximation accuracy with respect to the original geometry using fewer skeletal points.

Comparison with Point2Skeleton [LLL*21] More recently, using deep neural networks to predict skeletal representations from point clouds is beginning to be studied. Point2Skeleton [LLL*21] is a representative approach that directly learns skeletal representations from point clouds based on the MAT in an unsupervised manner. We use their pre-trained network on a large quantity of data for comparison.

Figure 7 as well as Table 2 show the comparison results. Note that Guitar and Chair-2 are from ShapeNet [CFG*15] dataset which is used for training Point2Skeleton. First, an obvious problem with Point2Skeleton is their limited generalization abilities; despite being trained on numerous data, it still generates unsatisfactory results for unseen shapes. Furthermore, it performs poorly on models with higher topology complexity, e.g., input with more than genus one. Also, the neural network shows incapability for capturing the properties of geometric transforms, such as the centrality of skeletal points. Conversely, our approach is robust to various shapes and can produce high-quality skeletal representations that have accurate geometries and faithful structures.

6. Discussions

In this section, in order to provide more insights and have better understandings of the proposed algorithm, we further give an in-depth analysis of some interesting properties of the Coverage Axis.

Table 2: Quantitative comparison on shape approximation errors among Point2Skeleton (P2S), Deep Point Consolidation (DPC) and Coverage Axis.

Model	P2S		DPC		Coverage Axis	
	$ V $	$\overleftarrow{\epsilon}$	$ V $	$\overleftarrow{\epsilon}$	$ V $	$\overleftarrow{\epsilon}$
Ant-2	100	16.412%	1194	8.863%	58	2.350%
Bottle	100	2.955%	1194	2.752%	14	2.956%
Chair-2	100	6.552%	1194	4.807%	89	2.890%
Dog	100	6.047%	1194	5.237%	49	2.174%
Dolphin	100	5.925%	1194	7.916%	49	1.971%
Fertility	100	8.162%	1194	4.226%	79	3.428%
Guitar	100	2.052%	1194	3.231%	60	2.032%
Hand-1	100	9.672%	1194	4.110%	44	3.441%
Human-2	100	6.517%	1194	4.758%	44	1.667%
Kitten	100	8.724%	1194	5.332%	50	3.450%
Snake	100	15.021%	1194	1.736%	40	1.309%
Average	-	8.003%	-	4.815%	-	2.515%

$|V|$ The number of skeleton points.

$\overleftarrow{\epsilon}$ Two-sided HD between original surface and reconstruction.

6.1. Centrality and Random Initialization

A good skeleton requires each skeletal point to be located at the center of its corresponding local geometry. Our method is able to effectively select the local centers for generating the skeletal representation. This property also enables our algorithm to handle randomly generated inner points instead of relying on computing the strictly defined Voronoi diagram as an input. In this section, we elaborate on this centrality property of our point selection strategy. We first generate the point candidates by *randomly* sampling inside




Figure 9: Skeletal point selection results from randomly sampled inner candidates. The blue points are candidate points generated by randomly sampling, while the orange points are selected points by our algorithm.




Figure 10: Effect of different numbers of randomly generated inner points on the centrality of the selected points. The blue points are randomly generated inner points while red points are selected skeletal points.

the shape and then applying our point selection algorithm introduced in Sec 4.2. The number of randomly generated inner points is 10000, and the offset is set to 0.02. A selection of results from these randomly generated points are shown in Figure 9. Thanks to the idea of optimal set coverage, it can be observed that the selected points are located at the centers of local geometry, even if the initial points are derived by random sampling. An evaluation of the influence of different randomly generated candidate inner point numbers $|P|$ for centrality is given in Figure 10. We find different numbers of candidate inner points do not show a large influence on the centrality.

In the following, we give an analysis of the rationale of the centrality property of our algorithm. Recall in Eq. 2, the constraint of the optimization is that the union of dilated balls needs to cover the sampled surface points. As explained in Sec. 6.4, the dilated radius should not deviate too much from the original radius, which means $r' \approx r$. Consider a local shape shown in Figure 11. The coverage goal of one candidate point P_a is defined by its farthest distance to the local geometry, e.g., d_a , given that all the surface samples are covered only if the farthest sample is covered. According to the definition of inner balls, the coverage ability is determined by its radius r_a (note that we have $r'_a \approx r_a$), which is actually the distance to the nearest surface samples of the local geometry.

The goal of finding the minimum number of dilated balls to cover all surface samples can be viewed as selecting those points with the most powerful coverage ability (cover all the target samples). That to say, its radius r_a should be as equal as possible to the distance to farthest point d_a so that we have $r'_a \geq d_a$ which satisfies the cover-




Figure 11: Illustration of the coverage of local geometry using medial balls. The orange points are surface samples. Balls centered at P_a and P_b are two candidate inner balls. The dilated balls are plotted using dashed circle.

age constraint. Those balls whose radii are close to the distance to the farthest point that needs to be covered will be preferred during optimization. Note that $r_a \approx d_a$ is consistent with the definition of the medial axis, which is intrinsically a set of centers of maximally inscribed spheres. In this case, the inner point P_b is preferred for the local geometry coverage since $d_b - r_b < d_a - r_a$. Note that we always have $d_b \geq r_b$. Finally, as shown in Figure 11, P_b is a better approximation than P_a of skeleton point and will be selected by our strategy.

6.2. Robustness to Noise

We evaluate the robustness of our method to surface noise. We compare with Q-MAT [LWS*15], which has the state-of-the-art performance in terms of approximation accuracy and goodness of structure (i.e., compactness and shape awareness). The number of vertices of the skeletal points are set the same for both methods. As shown in Figure 12, given noisy input, our method can still preserve the main features and achieve relatively high-precision shape approximation with a compact representation. The robustness mainly owes to the ball dilation strategy as well as the formulation based on global coverage. The dilated balls cover the insignificant details cause by local perturbations, while the overall coverage formulation gives a high-level optimization rather than focusing on local details. In contrast, Q-MAT [LWS*15], which uses step-by-step edge collapse based on local errors, cannot capture the fine details when the input surface has great noise.

6.3. Dilation Strategy

As aforementioned, the relationship between inner points and surface is built by the coverage, which is achieved by dilating the medial radius. There are two typical ways for inner ball dilation. One is by adding a specified value δ_r to the radius of each ball. The other is to scale the ball by multiplying the radius by a factor σ_r . We now discuss the effect of these two dilation strategies. To facilitate the discussion, we refer to the first way as *offset* and the second as *scaling*.

We take the model Octopus and Ant-3 as examples. Let r and r' denote the original medial radius and dilated medial radius respectively. For the offset manner, we set $r' = r + \delta_r$ where $\delta_r = 0.02$, while for the scaling manner, we set $r' = r \times \sigma_r$ with $\sigma_r = 1.5$.

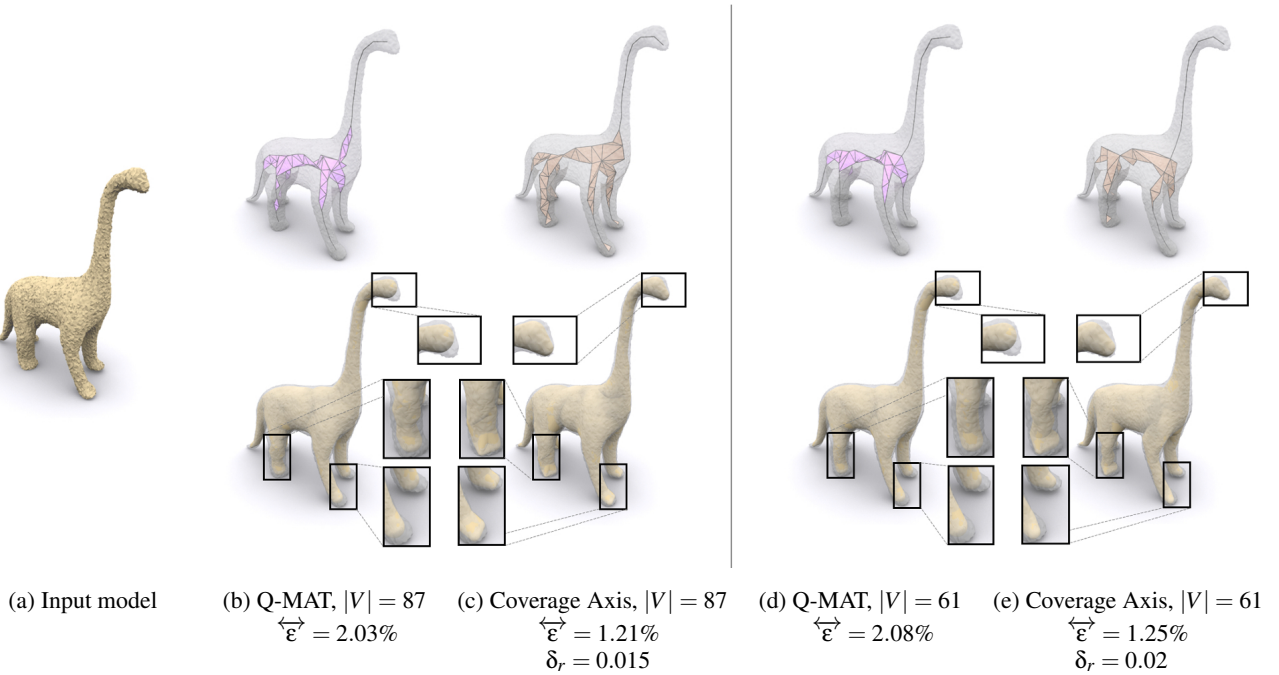


Figure 12: Comparison with Q-MAT on the robustness to surface noise. (a) Input model. (b)(c) and (d)(e) correspond to comparisons at different simplification levels.




Figure 13: Effect on different ball dilation strategies: (a) offset, (b) scaling.

As shown in Figure 13, the offset manner tends to generate selected points that are distributed more evenly, while the scaling manner can distinguish parts that have different radii, which leads to relatively large gap around the joints of different parts (as shown in the box region). The reason behind this is, adding a constant value is independent of the original radius, but the scaling makes the larger ball dilate more and control more area, which results in part differentiation.

6.4. Parameter Analysis

Dilation Factor The dilation factor δ_r is an important parameter that controls the simplicity of the output skeleton. In the following, we discuss the effect of different values of the parameter δ_r . We use

the Femur model as an example for demonstration, whose results are shown in Figure 14.

Obviously, a larger radius enables a ball to cover more surface points, which allows the algorithm to use fewer inner balls to cover the whole geometry. Therefore, a larger δ_r will lead to a higher simplified skeletal representation that captures less geometric details. Actually, the premise that we use surface information for simplification is that the envelope formed by dilated balls should not deviate significantly from the geometry of the input shape. On the one hand, we recommend not setting an excessively large δ_r , as this leads to the over-dilated balls of which union are inconsistent with the original shape structure. On the other hand, δ_r being too small is also discouraged because it runs counter to the simplicity and compactness, which are related requirements for a skeletal representation. In this paper, we always set $\delta_r = 0.02$ for our experiments.

Surface points and inner points. To investigate the influence of the number of surface points $|S|$ and candidate inner points $|P|$ (note it is determined by $|C|$, i.e., the number of samples to compute Voronoi diagram) on the algorithm performance, we conduct comprehensive experiments. We find that these two parameters do not show a significant influence on the selected point number $|V|$ and the quality of the resulting medial mesh. See more details in Appendix C. However, it has more effect on the time efficiency, for which a detailed analysis is given in Sec. 6.5.

6.5. Time Efficiency

We further discuss the time efficiency of our point selection strategy. The statistics is summarized in Table 3. We use Mixed-

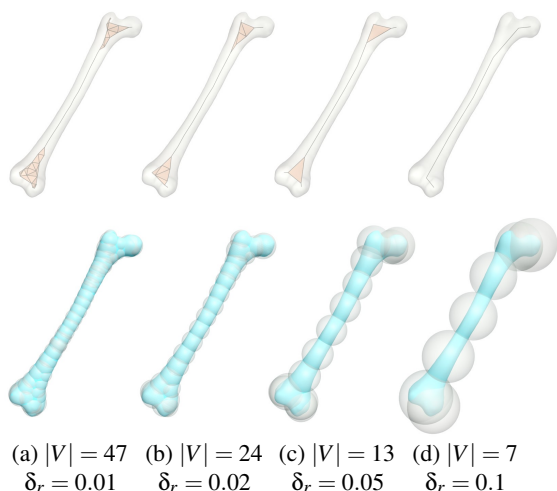


Figure 14: Simplification results using different dilation offsets. The first row and second row are simplification results and selected dilated coverage balls, respectively. We use $|V|$ to denote the number of vertices in the result and δ_r is the offset added to the original radius.

Table 3: Evaluation of the running time (s) of inner point selection.

Model	$ P $	$ V $	Time (s)	Model	$ P $	$ V $	Time (s)
Ant-1	13490	46	2.8	Hand-1	15475	42	6.7
Armadillo	12348	95	5.4	Hand-2	13665	60	48.8
Kitten	13718	45	36.2	Human-1	13023	45	0.9
Bunny	14837	106	10.1	Horse	13135	67	9.3
Crab	13971	58	3.3	Octopus-1	13432	49	0.3
Camel	12807	79	4.9	Octopus-2	13315	68	0.4
Dog	12623	49	13.1	Pig	13238	56	12.1
Dolphin	13998	43	55.4	Pliers	13097	31	0.9
Duck	17292	25	42.3	Snake	14160	39	0.2
Elephant	12503	87	7.8	Spider	13012	62	15.5
Eight	16393	28	16.1	Spectacle	11235	45	19.2
Fish	13279	53	42.1	Bear	13466	32	23.7
Fertility	13486	79	34.5	Vase	13174	106	35.2
Femur	14018	24	2.1	Venus	14362	36	9.7
Average							16.4

integer linear programming (MILP) solver in MATLAB (MathWorks 2021) on the platform explained in Sec 5.

As aforementioned, the experiments are conducted with mesh input using 1500 surface point samples, candidate inner points generated by 4000 surface samples. The dilation offset is set to 0.02. It can be viewed from Table 3 that the point selection only takes a few seconds for most shapes. For some cases, the complexity of the searching space of the Set Coverage problem may increase, which leads to one to two minutes of running time. The average running time for all the samples is 16.4 s.

Table 3 also reveals the difference in running time for various shapes, where the difference is because the complexity of the solution space in SCP varies for different geometries. Generally, solving for the coverage of plate-like geometries is more complex than tube-like ones, since the plate-like models have more eligible candidates that satisfy the coverage constraints, which leads to a larger

Table 4: Running time (with resulting vertex number in brackets) comparison (in seconds) with different methods using mesh and point cloud inputs.

Model	Mesh Input			Point Cloud Input		
	Q-MAT	SAT	Ours	DPC	P2S	Ours
Ant-1	3.1 (46)	38.2 (20k)	5.1 (46)	6.1 (1094)	3.2 (100)	6.3 (46)
Bunny	24.6 (106)	63.7 (13k)	14.9 (106)	7.2 (1094)	3.3 (100)	18.4 (106)
Horse	35.1 (67)	52.8 (12k)	11.5 (67)	8.2 (1094)	3.4 (100)	12.4 (67)
Pig	20.5 (56)	67.1 (13k)	14.3 (56)	6.6 (1094)	3.2 (100)	14.1 (56)
Venus	12.7 (36)	57.5 (8k)	13.7 (36)	7.1 (1094)	3.2 (100)	15.1 (36)

search space during optimization. Since many shapes consist of both tube-like and plate-like structures, there is a variation in running times. We further report running time comparison in Table 4.

SCP is a known NP-Hard problem [Har82] and its running time complexity is $O(|P|^2 2^{|S|})$ where $|P|$ and $|S|$ are the number of candidate inner points and surface points, respectively. To investigate the influence of $|S|$ and $|P|$ on running time efficiency and the resulting number of selected skeletal points $|V|$, we conduct additional experiments of which results are summarized in Table 5 and Table 6 respectively. We adopt the Octopus-1 model for evaluation. We fix $|C| = 4000$ in Table 5 and $|S| = 1500$ in Table 6. From the results, both $|S|$ and $|P|$ do not show a significant influence on the selected vertex number. However, $|S|$ has a larger impact on the running time, while $|P|$ shows less effect, which is consistent with the aforementioned analysis of time complexity.

Table 5: Effect of different numbers of surface samples $|S|$ on selected point number and running time efficiency.

$ S $	$ V $	Time (s)
500	47	0.25
1500	49	0.29
2500	49	0.34
3500	49	0.73
4500	49	0.98
5500	49	1.33
6500	49	2.44
7500	49	3.24
8500	49	4.22

Table 6: Effect of different numbers of inner point candidates $|P|$ on selected point number and running time efficiency. C denotes surface samples used for generating candidate inner points.

$ C $	$ P $	$ V $	Time (s)
1000	2964	48	0.08
2000	6894	49	0.18
3000	10776	49	0.28
4000	13432	49	0.31
5000	16749	49	0.49
6000	22301	49	0.82
7000	26161	49	1.23
8000	29859	49	1.43
9000	33651	49	1.65

To evaluate the time efficiency for different algorithm settings, we conduct an additional experiment where the initial candidate points are generated by random sampling. Here, we randomly sample 10000 points inside a shape and test the time efficiency using two different dilation strategies and different dilation factors. The results are given in Table 7. It can be observed the running time is stable in different configurations.

7. Additional Features of Coverage Axis

In this section, we analyze some of the other interesting features of our method, including the ability to handle large models and the flexibility of user-specified point control.

Table 7: Additional evaluations on time efficiency. We show the running time (s) of point selection from randomly generated candidates using different ball dilation strategies.

Model	Ant-2				FEMUR				Hand-2				Fertility			
Offset	0.015	0.02	0.025	0.03	0.015	0.02	0.025	0.03	0.015	0.02	0.025	0.03	0.015	0.02	0.025	0.03
Selected points	98	59	48	34	39	26	21	17	86	63	57	38	153	82	62	49
Time (s)	3.2	2.1	8.1	7.5	3.3	2.2	1.2	0.3	19.5	15.8	10.2	9.8	21.0	17.4	9.3	10.7
Scaling	1.5	1.75	2	2.25	1.5	1.75	2	2.25	1.5	1.75	2	2.25	1.5	1.75	2	2.25
Selected points	98	77	62	49	30	17	12	10	65	31	22	18	80	64	40	31
Time (s)	3.1	1.4	0.9	0.4	0.4	0.8	1.7	0.3	0.7	3.3	3.2	2.5	0.7	0.5	0.7	1.6




Figure 15: Coverage Axis with divide and conquer strategy. The model is divided into 29 parts for solving the set coverage problem separately on each part. The average running time of inner point selection of each component is 1.04 s. However, it takes more than 15 min to solve for the entire model.

7.1. Divide and Conquer Strategy for Large Models

Since the set coverage problem is NP-hard, the action space grows exponentially with respect to the number of surface points. Fortunately, an interesting feature of the proposed method is the flexibility in handling large complex models. The overall insights lie in the combination of divide and conquer strategy and our set coverage formulation. In this case, we take a mesh model with Voronoi initialization for explanation.

The first step is decomposing a mesh into smaller and meaningful sub-meshes with some existing methods, e.g., [YL21; HHF*19]. Then we have surface samples divided into several subsets S_1, S_2, \dots, S_n where n is the number of surface components. According to the correspondence between the vertices on the inside Voronoi diagram and the surface samples, we partition the inside Voronoi diagram into P_1, P_2, \dots, P_n accordingly. For each component i , with surface samples S_i and inner point candidates P_i , we solve the SCP and obtain the selection result P_i^+ embedded on the sub-Voronoi diagram. This process reduces the size of the problem and speeds up our algorithm by divide and conquer.

Finally, we merge each Voronoi diagram as well as the selected points. Then, we have all the selected points embedded on the original inside Voronoi diagram so that the model's simplification can be achieved by removing redundant components as we described in Sec. 4.3. We demonstrate the results in Figure 15 (a-b). Note that the complex model is with 30000 surface samples which are further used for generating candidate inner points.

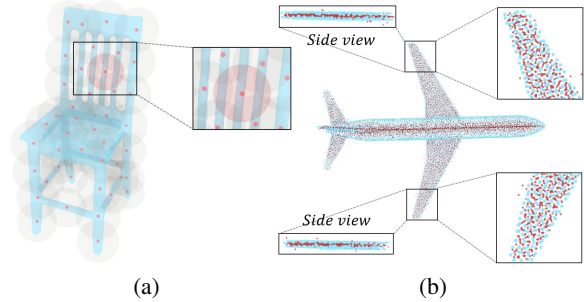


Figure 16: Examples of failure case. (a) The over-dilated inner balls destroy the original shape topology. (b) Incorrect candidate initialization of thin flat model taken point clouds as the input. The red points are candidate points and the blue points are the input, where some candidates are located outside the shape.

7.2. User-specified Point Control

Our method allows users to specify the ignorable surface parts that can be neglected. To indicate the ignorable surface points, the constraint in Eq 2 are reformulated as $\mathbf{Dv} \geq \mathbf{B}$, where the corresponding values for the ignorable surface points in \mathbf{B} are set to 0 and the others to 1. In this way, the optimization tends to remove the skeletal points in the specified parts to favor parsimony. Users can also directly indicate if a skeletal point should be preserved or removed, for which we just simply set the corresponding values in \mathbf{v} to 0 or 1 after the point selection.

8. Limitations

In this section, we discuss the limitations of our method. First, the number of selected points is related to the dilation parameter δ_r , which leads to a difficulty in explicit control of the number of vertices of a skeleton. In particular, when inner balls are excessively dilated, the topology of the input surface tends to be destroyed by the coverage. See Figure 16 (a). Another limitation lies in the difficulty in preserving approximation accuracy when computing extremely decimated representations (lower vertices number in the skeleton). As shown in Figure 17, Q-MAT [LWS*15] shows a better control of the reconstruction error $\sqrt{\epsilon}$ than ours, as the number of target skeletal points $|V|$ decreases. This is because Coverage Axis relies on much larger dilation factors for computing highly decimated results, which brings difficulty in preserving the local geometry. Detailed statistics can be found in Appendix D.

Our method usually suffers from a relatively high computational

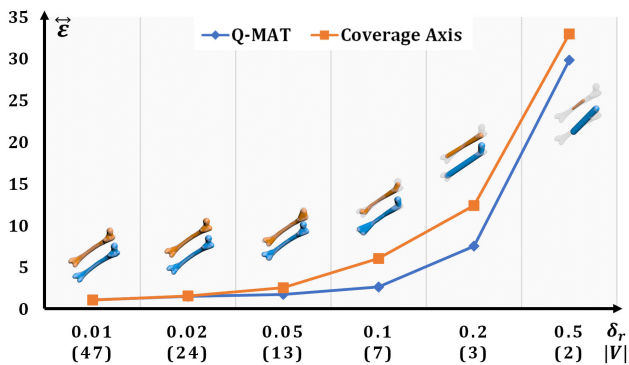


Figure 17: The change of the approximation error (two-sided HD ϵ) w.r.t. the simplification level. Compared with Q -MAT, our method exhibits larger errors for highly decimated medial surfaces. We adopt the FEMUR model for evaluation.

cost when given a large number of surface points to be covered. To alleviate this, we demonstrate a potential approach in designing a divide and conquer strategy based on geometric clues to achieve acceleration. In addition, for the point cloud input, especially the thin flat models, it remains an open problem to clearly distinguish between inner and outside volume. Therefore, for this kind of input, it is extremely challenging to generate legitimate candidate points that are located inside the shape, which can result in failure cases during point selection (Figure 16 (b)).

9. Conclusion

In this paper, we propose a novel, simple, yet effective formulation named Coverage Axis for 3D shape skeletonization of both meshes and point clouds. Inspired by the set cover problem (SCP), our goal is to cover all the surface points using as few inside medial balls as possible. Compared with previous methods, our approach shows a set of appealing properties including its simple formulation, high-level abstraction, robustness to noise, flexibility to handle various cases, and so on. Comprehensive experiments verify the aforementioned claims. In the future, we believe Coverage Axis, as a tool for encoding shapes, has a large potential value in various applications including shape abstraction, shape segmentation, simulation as well as animation, etc. Another point worth noting is the combination of Coverage Axis and learning techniques to improve the work of shape analysis and processing, such as 3D object detection, shape classification, automatic blending and so on.

References

[ABE09] ATTALI, DOMINIQUE, BOISSONNAT, JEAN-DANIEL, and EDELSBRUNNER, HERBERT. “Stability and computation of medial axes—a state-of-the-art report”. *Mathematical foundations of scientific visualization, computer graphics, and massive data exploration*. Springer, 2009, 109–125 2.

[ACK01] AMENTA, NINA, CHOI, SUNGHEE, and KOLLURI, RAVI KRISHNA. “The power crust”. *Proceedings of the sixth ACM symposium on Solid modeling and applications*. 2001, 249–266 1–3, 5.

[AM96] ATTALI, DOMINIQUE and MONTANVERT, ANNICK. “Modeling noise for a better simplification of skeletons”. *Proceedings of 3rd IEEE International Conference on Image Processing*. Vol. 3. IEEE. 1996, 13–16 2.

[ATC*08] AU, OSCAR KIN-CHUNG, TAI, CHIEW-LAN, CHU, HUNG-KUO, et al. “Skeleton extraction by mesh contraction”. *ACM transactions on graphics (TOG)* 27.3 (2008), 1–10 1, 2.

[Aur87] AURENHAMMER, F. “Power Diagrams: Properties, Algorithms and Applications”. *SIAM J. Comput.* 16.1 (Feb. 1987), 78–96. ISSN: 0097-5397. DOI: 10.1137/0216006. URL: <https://doi.org/10.1137/0216006> 3.

[BDS*18] BARILL, GAVIN, DICKSON, NEIL, SCHMIDT, RYAN, et al. “Fast Winding Numbers for Soups and Clouds”. *ACM Transactions on Graphics* (2018) 4.

[Blu*67] BLUM, HARRY et al. *A transformation for extracting new descriptors of shape*. Vol. 43. MIT press Cambridge, MA, 1967 1, 2.

[BP07] BARAN, ILYA and POPOVIĆ, JOVAN. “Automatic rigging and animation of 3d characters”. *ACM Transactions on graphics (TOG)* 26.3 (2007), 72–es 1.

[CCT11] CHAUSSARD, JOHN, COUPRIE, MICHEL, and TALBOT, HUGUES. “Robust skeletonization using the discrete λ -medial axis”. *Pattern Recognition Letters* 32.9 (2011), 1384–1394 2.

[CFG*15] CHANG, ANGEL X, FUNKHOUSER, THOMAS, GUIBAS, LEONIDAS, et al. “Shapenet: An information-rich 3d model repository”. *arXiv preprint arXiv:1512.03012* (2015) 8.

[CL05] CHAZAL, FRÉDÉRIC and LIEUTIER, ANDRÉ. “The “ λ -medial axis””. *Graphical Models* 67.4 (2005), 304–331 2.

[CTO*10] CAO, JUNJIE, TAGLIASACCHI, ANDREA, OLSON, MATT, et al. “Point cloud skeletons via laplacian based contraction”. *2010 Shape Modeling International Conference*. IEEE. 2010, 187–197 3.

[CZC*20] CHENG, JINGLIANG, ZHENG, XINYU, CHEN, SHUANGMIN, et al. “Skeletonization via dual of shape segmentation”. *Computer Aided Geometric Design* 80 (2020), 101856 2.

[DEGN98] DEY, TAMAL K, EDELSBRUNNER, HERBERT, GUHA, SUMANTA, and NEKHAYEV, DMITRY V. “Topology preserving edge contraction”. *Publ. Inst. Math.(Beograd)(NS. Citeseer*. 1998 5.

[DS06] DEY, TAMAL K and SUN, JIAN. “Defining and computing curve-skeletons with medial geodesic function”. *Symposium on geometry processing*. Vol. 6. 2006, 143–152 1.

[DXX*20] DOU, ZHIYANG, XIN, SHIQING, XU, RUI, et al. “Top-down shape abstraction based on greedy pole selection”. *IEEE Transactions on Visualization and Computer Graphics* (2020) 1.

[DZ02] DEY, TAMAL K and ZHAO, WULUE. “Approximate medial axis as a Voronoi subcomplex”. *Proceedings of the seventh ACM symposium on Solid modeling and applications*. 2002, 356–366 2.

[FLM03] FOSKEY, MARK, LIN, MING C, and MANOCHA, DINESH. “Efficient computation of a simplified medial axis”. *J. Comput. Inf. Sci. Eng.* 3.4 (2003), 274–284 2.

[FTB13] FARAJ, NOURA, THIERY, JEAN-MARC, and BOUBEKEUR, TAMY. “Progressive medial axis filtration”. *SIGGRAPH Asia 2013 Technical Briefs*. 2013, 1–4 2.

[GH97] GARLAND, MICHAEL and HECKBERT, PAUL S. “Surface simplification using quadric error metrics”. *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. 1997, 209–216 2, 3, 5.

[GMPW09] GIESEN, JOACHIM, MIKLOS, BALINT, PAULY, MARK, and WORMSER, CAMILLE. “The scale axis transform”. *Proceedings of the twenty-fifth annual symposium on Computational geometry*. 2009, 106–115 2.

[Har82] HARTMANIS, JURIS. “Computers and intractability: a guide to the theory of np-completeness (michael r. Garey and david s. Johnson)”. *Siam Review* 24.1 (1982), 90 11.

[HHF*19] HANOCCA, RANA, HERTZ, AMIR, FISH, NOA, et al. “MeshCNN: A network with an edge”. *ACM Transactions on Graphics (TOG)* 38.4 (2019), 1–12 12.

[HWC*13] HUANG, HUI, WU, SHIHAO, COHEN-OR, DANIEL, et al. “L1-medial skeleton of point cloud.” *ACM Trans. Graph.* 32.4 (2013), 65–1 3.

- [HWG*13] HUANG, HUI, WU, SHIHAO, GONG, MINGLUN, et al. "Edge-aware point set resampling". *ACM transactions on graphics (TOG)* 32.1 (2013), 1–12 5.
- [LCLT07] LIPMAN, YARON, COHEN-OR, DANIEL, LEVIN, DAVID, and TAL-EZER, HILLEL. "Parameterization-free projection for geometry reconstruction". *ACM Transactions on Graphics (TOG)* 26.3 (2007), 22–es 3.
- [LGS12] LIVESU, MARCO, GUGGERI, FABIO, and SCATENI, RICCARDO. "Reconstructing the curve-skeletons of 3d shapes using the visual hull". *IEEE transactions on visualization and computer graphics* 18.11 (2012), 1891–1901 2, 3.
- [LLL*20] LIN, CHENG, LIU, LINGJIE, LI, CHANGJIAN, et al. "SEG-MAT: 3D shape segmentation using medial axis transform". *IEEE Transactions on Visualization and Computer Graphics* (2020) 1.
- [LLL*21] LIN, CHENG, LI, CHANGJIAN, LIU, YUAN, et al. "Point2Skeleton: Learning skeletal representations from point clouds". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021, 4277–4286 3, 7, 8.
- [LWS*15] LI, PAN, WANG, BIN, SUN, FENG, et al. "Q-MAT: Computing medial axis transform by quadratic error minimization". *ACM Transactions on Graphics (TOG)* 35.1 (2015), 1–16 2, 3, 5, 6, 9, 12.
- [MGP10] MIKLOS, BALINT, GIESEN, JOACHIM, and PAULY, MARK. "Discrete scale axis representations for 3D geometry". *ACM SIGGRAPH 2010 papers*. 2010, 1–10 2, 6, 7.
- [MLM16] MARIE, ROMAIN, LABBANI-IGBIDA, OUIDDAD, and MOUADDIB, EL MUSTAPHA. "The delta medial axis: a fast and robust algorithm for filtered skeleton extraction". *Pattern Recognition* 56 (2016), 26–39 3.
- [MLT88] MAGNENAT-THALMANN, NADIA, LAPERRIERE, RICHARD, and THALMANN, DANIEL. "Joint-dependent local deformations for hand animation and object grasping". In *Proceedings on Graphics interface '88*. Citeseer. 1988 7.
- [MWO03] MA, WAN-CHUN, WU, FU-CHE, and OUHYOUNG, MING. "Skeleton extraction of 3D objects with radial basis functions". *2003 Shape Modeling International*. IEEE. 2003, 207–215 1, 2.
- [NBPFI1] NATALI, MATTIA, BIASOTTI, SILVIA, PATANÈ, GIUSEPPE, and FALCIDIENO, BIANCA. "Graph-based representations of point clouds". *Graphical Models* 73.5 (2011), 151–164 2.
- [RAV*19] REBAIN, DANIEL, ANGLÉS, BAPTISTE, VALENTIN, JULIEN, et al. "LSMAT Least squares medial axis transform". *Computer Graphics Forum*. Vol. 38. 6. Wiley Online Library. 2019, 5–18 3.
- [RLS*21] REBAIN, DANIEL, LI, KE, SITZMANN, VINCENT, et al. "Deep medial fields". *arXiv preprint arXiv:2106.03804* (2021) 3.
- [SBH*11] SHEN, WEI, BAI, XIANG, HU, RONG, et al. "Skeleton growing and pruning with bending potential ratio". *Pattern Recognition* 44.2 (2011), 196–209 3.
- [SCYW13] SUN, FENG, CHOI, YI-KING, YU, YIZHOU, and WANG, WENPING. "Medial meshes for volume approximation". *arXiv preprint arXiv:1308.3917* (2013) 1, 3.
- [SFC*11] SHOTTON, JAMIE, FITZGIBBON, ANDREW, COOK, MAT, et al. "Real-time human pose recognition in parts from single depth images". *CVPR 2011*. Ieee. 2011, 1297–1304 1.
- [SFM07] SUD, AVNEESH, FOSKEY, MARK, and MANOCHA, DINESH. "Homotopy-preserving medial axis simplification". *International Journal of Computational Geometry & Applications* 17.05 (2007), 423–451 2.
- [SKS11] STOLPNER, SVETLANA, KRY, PAUL, and SIDDIQI, KALEEM. "Medial spheres for shape approximation". *IEEE transactions on pattern analysis and machine intelligence* 34.6 (2011), 1234–1240 1.
- [SLSK07] SHARF, ANDREI, LEWINER, THOMAS, SHAMIR, ARIEL, and KOBBELT, LEIF. "On-the-fly Curve-skeleton Computation for 3D Shapes". *Computer Graphics Forum*. Vol. 26. 3. Wiley Online Library. 2007, 323–328 2.
- [SSGD03] SUNDAR, HARI, SILVER, DEBORAH, GAGVANI, NIKHIL, and DICKINSON, SVEN. "Skeleton based shape matching and retrieval". *2003 Shape Modeling International*. IEEE. 2003, 130–139 7.
- [TAOZ12] TAGLIASACCHI, ANDREA, ALHASHIM, IBRAHEEM, OLSON, MATT, and ZHANG, HAO. "Mean curvature skeletons". *Computer Graphics Forum*. Vol. 31. 5. Wiley Online Library. 2012, 1735–1744 1–3.
- [TDS*16] TAGLIASACCHI, ANDREA, DELAME, THOMAS, SPAGNUOLO, MICHELA, et al. "3d skeletons: A state-of-the-art report". *Computer Graphics Forum*. Vol. 35. 2. Wiley Online Library. 2016, 573–597 3.
- [TGB13] THIERY, JEAN-MARC, GUY, ÉMILIE, and BOUBEKEUR, TAMY. "Sphere-meshes: Shape approximation using spherical quadric error metrics". *ACM Transactions on Graphics (TOG)* 32.6 (2013), 1–12 2.
- [TGB16] THIERY, JEAN-MARC, GUY, ÉMILIE, BOUBEKEUR, TAMY, and EISEMANN, ELMAR. "Animated mesh approximation with sphere-meshes". *ACM Transactions on Graphics (TOG)* 35.3 (2016), 1–13 7.
- [THP*19] TANG, JIAPENG, HAN, XIAO GUANG, PAN, JUNYI, et al. "A skeleton-bridged deep learning approach for generating meshes of complex topologies from single rgb images". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, 4541–4550 1.
- [TZC09] TAGLIASACCHI, ANDREA, ZHANG, HAO, and COHEN-OR, DANIEL. "Curve Skeleton Extraction from Incomplete Point Cloud". *ACM Trans. Graph.* 28.3 (July 2009). ISSN: 0730-0301. DOI: 10.1145/1531326.1531377 3.
- [WHG*15] WU, SHIHAO, HUANG, HUI, GONG, MINGLUN, et al. "Deep points consolidation". *ACM Transactions on Graphics (ToG)* 34.6 (2015), 1–13 1, 3, 7.
- [XZKS19] XU, ZHAN, ZHOU, YANG, KALOGERAKIS, EVANGELOS, and SINGH, KARAN. "Predicting animation skeletons for 3d articulated models via volumetric nets". *2019 International Conference on 3D Vision (3DV)*. IEEE. 2019, 298–307 1, 2.
- [YL21] YAZ, ILKER O. and LORIOT, SÉBASTIEN. "Triangulated Surface Mesh Segmentation". *CGAL User and Reference Manual*. 5.3. CGAL Editorial Board, 2021. URL: <https://doc.cgal.org/5.3/Manual/packages.html#PkgSurfaceMeshSegmentation> 12.
- [YLJ18] YAN, YAJIE, LETSCHER, DAVID, and JU, TAO. "Voxel Cores: Efficient, robust, and provably good approximation of 3d medial axes". *ACM Transactions on Graphics (TOG)* 37.4 (2018), 1–13 3.
- [YRL*21] YANG, YIDING, REN, ZHOU, LI, HAOXIANG, et al. "Learning dynamics via graph neural networks for human pose estimation and tracking". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021, 8074–8084 1.
- [YSC*16] YAN, YAJIE, SYKES, KYLE, CHAMBERS, ERIN, et al. "Erosion thickness on medial axes of 3D shapes". *ACM Transactions on Graphics (TOG)* 35.4 (2016), 1–12 3.
- [YYG18] YANG, BAORONG, YAO, JUNFENG, and GUO, XIAOHU. "DMAT: Deformable medial axis transform for animated mesh approximation". *Computer Graphics Forum*. Vol. 37. 7. Wiley Online Library. 2018, 301–311 1, 7.
- [YYW*20] YANG, BAORONG, YAO, JUNFENG, WANG, BIN, et al. "P2MAT-Net: Learning medial axis transform from sparse point clouds". *Computer Aided Geometric Design* 80 (2020), 101874 3.