











Coverage Axis++: Efficient Inner Point Selection for 3D Shape Skeletonization

Zimeng Wang*¹ , Zhiyang Dou*^{1,2} , Rui Xu³ , Cheng Lin¹ , Yuan Liu¹ , Xiaoxiao Long¹ ,
Shiqing Xin³ , Taku Komura¹ , Xiaoming Yuan^{†1} , Wenping Wang^{†4} 

¹The University of Hong Kong ²TransGP ³Shandong University ⁴Texas A&M University

*Equal Contribution; †Equal Advising.




Figure 1: Coverage Axis++ enables highly efficient skeletonization of the given shapes represented by water-tight meshes, triangle soups or point clouds. Compared with the SOTA method [DLX*22; LWS*15; PWG*19; LLL*21; CD23], Coverage Axis++ allows for fast computation of compact medial axis approximation with low reconstruction error; i.e., Coverage Axis++: 5.362% reconstruction error in 2.9s v.s. Coverage Axis: 6.574% reconstruction error in 236.1s on average; See Tab. 1.

Abstract

We introduce Coverage Axis++, a novel and efficient approach to 3D shape skeletonization. The current state-of-the-art approaches for this task often rely on the watertightness of the input [LWS*15; PWG*19; PWG*19] or suffer from substantial computational costs [DLX*22; CD23], thereby limiting their practicality. To address this challenge, Coverage Axis++ proposes a heuristic algorithm to select skeletal points, offering a high-accuracy approximation of the Medial Axis Transform (MAT) while significantly mitigating computational intensity for various shape representations. We introduce a simple yet effective strategy that considers shape coverage, uniformity, and centrality to derive skeletal points. The selection procedure enforces consistency with the shape structure while favoring the dominant medial balls, which thus introduces a compact underlying shape representation in terms of MAT. As a result, Coverage Axis++ allows for skeletonization for various shape representations (e.g., water-tight meshes, triangle soups, point clouds), specification of the number of skeletal points, few hyperparameters, and highly efficient computation with improved reconstruction accuracy. Extensive experiments across a wide range of 3D shapes validate the efficiency and effectiveness of Coverage Axis++. Our codes are available at https://github.com/Frank-ZY-Dou/Coverage_Axis.

CCS Concepts

• Computing methodologies → Shape analysis;

1. Introduction

Skeletal representations have become a popular tool in various applications of shape analysis and geometric processing, as they efficiently capture the underlying structures of 3D shapes. They have been widely adopted for various tasks including

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; DWL*23; LXC*21], and animation [BP07; YYG18], among others. With the development of deep learning, skeletal representations also facilitate learning-based approaches in various

tasks [HWQ*19; RLS*21; THT*21; HCY*22; PGT*24; WYT23; XZX*24; WTJL24; GWHM24].

Previous efforts have been made for the computation of curve skeleton [ATC*08; MWO03; TAOZ12; XZKS19], which consists of only 1D curves. Dey and Sun [DS06] provide a mathematical formulation based on the Medial Geodesic Function, yet curve skeletons are predominantly comprehended empirically.

Medial Axis Transform (MAT) [Blu*67], as a complete shape descriptor, is another popular skeletal representation. The MAT is defined by a union of maximally inscribed balls within the shape, accompanied by their respective radius functions. A formal definition of MAT can be found in Sec. 3.1. In contrast to curve skeletons, the MAT encompasses both curve-like and surface-like structures, offering a consistent definition for arbitrary shapes and superior representational capability.

The MAT computation poses challenges due to sensitivity to boundary noise [LWS*15] and strict input geometry requirements, e.g., watertightness and manifoldness of the surface [LWS*15; PWG*19; WWWG22]. To tackle the problem, Coverage Axis [DLX*22] is introduced to model skeletal point selection as a Set Cover Problem (SCP), which does not rely on surface connectivity and is able to handle both mesh or point cloud inputs. It aims to identify the smallest sub-collection covering the entire shape. This method minimizes dilated inner balls to approximate the shape, leading to a compact representation of the original shape. However, solving the SCP problem, which is known to be an NP-hard problem, can be time-consuming in some cases. For instance, the average running time of the Coverage Axis can be 236.1s with the standard deviation being 408.4s; See Tab. 1. Meanwhile, the recent work Neural Skeleton [CD23] proposes to first learn a Signed Distance Field from the given shape and then use the Coverage Axis to get the final compact skeleton of the given mesh or point cloud. However, it still suffers from low runtime performance (59.27s compared to our 1.63s) and relatively low reconstruction accuracy as shown in Tab. 3.

In this paper, we develop a simple yet more efficient skeletonization method for shapes represented as meshes or point clouds. We take shape representation capability, centrality, and uniformity into consideration and develop a heuristic algorithm to achieve representative skeletal point selection with high efficiency. Specifically, we present a scoring scheme that assigns a score to each inner point as a quantification of its skeletal representation ability. The scoring consists of the *coverage*, *uniformity* and *centrality* scores. The coverage score is determined by evaluating the number of surface samples that can be covered within the dilated ball associated with the candidate to ensure coverage inspired by [DLX*22]. The coverage score ensures consistency with the shape structure by favoring interior points that dominate a larger area, aligning with the definition of the Medial Axis Transform (MAT) as the set of centers of maximally inscribed spheres. On the other hand, uniformity is one of the key desired features of compact skeletonization to suppress redundant structures, for which we introduce the uniformity score. The score is calculated by measuring the distance between the candidate and the nearest point in the set of selected inner points. This achieves a uniformly distributed skeletonization result by preventing the selected points from clustering together. More-

over, the centrality score is employed to encourage those central points to receive higher scores to promote centrality further. The skeletal points are derived by running a priority queue defined by an integration of these scores. We show that this simple yet effective point selection manner produces a better abstraction of the overall shape as well as a shape-aware point distribution without the need to solve an optimization problem with high computation complexity. Notably, although the priority queue works in a greedy manner, Coverage Axis++ achieves global consistency between the skeleton and the original shape and allows for specifying different numbers of target skeletal points; See Sec. 6.1.

Compared with the existing methods [ACK01; LWS*15; GMPW09; FLM03; DZ02; SFM07; MGP10; DLX*22], Coverage Axis++ surpasses all existing methods in terms of both MAT approximation accuracy and computation efficiency while being able to handle various inputs, including mesh, point cloud or polygon soup. To summarize, our method allows for 1) skeletonization for various shape representations including water-tight meshes, triangle soups and point clouds; 2) specifying the number of skeletal points, which is a desirable feature for learning-based methods as a fixed number of points are typically required [HWQ*19; THT*21; HCY*22; PGT*24]; 3) highly efficient computation while achieving better or more competitive reconstruction accuracy; 4) few hyperparameters, i.e., only dilation factor is required; 5) randomly generated candidates inside the volume as Coverage Axis++ selects the most expressive ones from overfilled inner point candidates.

We conduct extensive experiments to demonstrate the effectiveness and robustness of our method on a variety of 3D shapes. The comprehensive evaluation reveals that our method effectively captures the underlying MATs with lower shape approximation errors while achieving much higher computation efficiency.

2. Related Work

Efficient shape representation is essential for various applications of geometric modeling, shape analysis and shape generation [MSLJ23; LGC*23; SKSJ20; XWD*22; XDW*23; DXX*20; FXX*22; WYT23; LLL*20; LYZ*23; YLL*23; ZLY*23; YPN*23; YDL*23; Ale21; PGT*24]. Skeletonization, as a popular choice in this regard, has been extensively studied in recent decades. We refer readers to [TDS*16] for a detailed survey covering various forms of skeletons. In the following, we mainly review the related works in Medial Axis Transform (MAT) computation for mesh and point cloud inputs.

Medial Axis Transform

Mesh Inputs. Efforts in MAT computation given the watertight mesh inputs have been extensive [ACK01; GWM23; FLM03; DZ02; SFM07; HWY*23; GYY*23; HCY*22; DLX*22; RLS*21; WWWG22; LLL*24; LFWN20]. Specifically, λ -medial axis [CL05] define weak feature size as a pruning criterion but struggle with feature preservation at different scales [ABE09]. Scale Axis Transform [MGP10] excels in spike pruning but is limited by high computational cost and topology disruption. Progressive medial axis [FTB13] proposes to perform topology-preserved edge collapse based on sphere absorption. Many studies have been conducted for MAT simplification using different

metrics [MLM16; SBH*11; YSC*16; YLJ18; SCYW13; LWS*15; PWG*19]. Li et al. [LWS*15] propose Q-MAT for highly accurate shape approximation, utilizing a quadratic error metric [GH97] during simplification. Despite relatively high approximation accuracy, Q-MAT relies on watertight surfaces. It is sensitive to the quality of MAT initialization, as revealed by [DLX*22]. Later, Q-MAT+ [PWG*19], as a variant of Q-MAT, further introduces more mesh information like shape diameter function to improve the approximation accuracy and robustness of the original Q-MAT method, especially for those sharp features. Recently, MATFP [WWG22] is developed to compute MAT with well-preserved geometric features via the Restricted Power Diagram for CAD mesh models. Wang et al. [WHS*24] have also explored addressing the topology preservation issue by employing Restricted Power Diagrams.


Point Cloud Inputs. Point cloud skeletonization is still a fascinating yet challenging problem [And09; CTO*10; LGS12; LZG*23; HWC*13; RAV*19; WHG*15; YYW*20; DLX*22] due to the absence of manifold information regarding the underlying surface of the shape. Specifically, L1-medial skeleton [HWC*13] employs Locally Optimal Projection [LCLT07] to contract point clouds to form the skeleton. LSMAT [RAV*19] approximates the MAT based on the Signed Distance Function (SDF) from a densely sampled oriented point set. Wu et al. [WHG*15] considers linking surface points to the skeletal points on the meso-skeleton [TAOZ12]. Yet, this representation only yields unstructured point sets, which lacks topological constraints describing the skeleton connectivity and leads to low approximation accuracy. Coverage Axis [DLX*22] models skeletal point selection as a Set Cover Problem (SCP), which does not rely on surface connectivity, thus allowing for point cloud inputs. However, it suffers from high computational costs as SCP is NP-hard, especially when facing shapes with planar structures (See detailed analysis in Appendix D).

With the advent of deep learning, learning-based approaches have been explored for point clouds skeletonization [YYW*20; LLL*21; GYY*23]. For example, Point2Skeleton [LLL*21] learns to predict a set of skeletal points by learning a geometric transformation and then analyzes the connectivity of the skeletal points to form mesh structures from point clouds. However, it suffers from limited generalization due to dependency on training data and lacks accuracy for geometric features. Recently, Neural Skeleton [CD23] leverages Implicit Neural Representations (INR), i.e., SDF, together with coverage axis [DLX*22] for skeleton computation, yet it still faces challenges with long running time and dependence on consistently oriented normals for input shapes.

3. Preliminaries

3.1. Medial Axis Transform

For a closed, oriented, and bounded two-manifold surface \mathcal{S} in \mathbb{R}^3 , the *medial axis* is defined as the locus of the centers of *maximally inscribed spheres* (medial balls) that are tangent to \mathcal{S} at two or more points. Denoted by a pair $(\mathcal{M}, \mathcal{R})$, where \mathcal{M} is the medial axis of \mathcal{S} and \mathcal{R} is its radius function, the combination forms the *Medial Axis Transform* (MAT).



3.2. Voronoi Diagram & Power Diagram

Voronoi Diagram (VD) 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 VD inside a shape generated by surface samples [ACK01; LWS*15; DLX*22].

In Power Diagram (PD) [Aur87], 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

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

Therefore, a generator with a larger weight is more dominant.

4. Method

Our goal is to select skeletal points from a set of candidate points denoted as $P = \{p_i\}$, by utilizing a group of sampled points represented as $S = \{s_j\}$ on the surface, where each point is defined by its 3D Cartesian coordinates in \mathbb{R}^3 . We present an effective and efficient algorithm for skeletal point selection. We denote by P^+ the set of selected skeletal points and P^- the set of unselected ones, such that $P = P^+ \cup P^-$ always holds. We gradually select a point from P^- and add it to P^+ until the selected points reach the number specified by the user. In this paper, all inner point candidates are randomly generated inside the volume; See an example in Fig. 2 (b). The algorithm for shape skeletonization is outlined through the following steps: scoring each candidate point for point selection and establishing connectivity among skeletal points.

4.1. Coverage Score

To achieve compact skeletonization, we aim to select representative points that capture the key features of the given 3D object by measuring its shape coverage [DLX*22] capability.

We first compute the medial radii $R = \{r_i\}$ for the candidate points, where r_i is defined as the closest distance of p_i to the sampled surface set S , i.e.,

$$r_i := \text{dist}(p_i, S) = \min_{s_j \in S} \|p_i - s_j\|. \quad (1)$$

The set R together with the center points P define the candidate balls. We slightly dilate all the balls by a small factor to their radii by offset, namely,

$$r'_i = r_i + \delta_r, \quad (2)$$

leading to a set of dilated balls $B = \{B_i\}$, where B_i is the ball centered at p_i with radius r'_i . A matrix $\mathbf{D} \in \{0, 1\}^{m \times n}$ is introduced by [DLX*22], where m and n are total numbers of sampled surface points and candidate skeletal points, respectively. Each entry $d_{ji} \in \{0, 1\}$ of \mathbf{D} indicates if the surface point s_j is covered by the dilated ball (p_i, r'_i) :

$$d_{ji} = \begin{cases} 1, & \text{if } \|p_i - s_j\|_2 \leq r'_i, j = 1, \dots, m, i = 1, \dots, n \\ 0, & \text{if } \|p_i - s_j\|_2 > r'_i, j = 1, \dots, m, i = 1, \dots, n. \end{cases} \quad (3)$$




Figure 2: The pipeline of Coverage Axis++. (a) Input point cloud. (b) Inner point candidates. (c) Candidate inner balls with original radius. (d) Selected inner balls with original radius. (e) Selected inner points. (f) Skeletonization result.

According to Eq. 3, we have the coverage matrix $\mathbf{D} = (d_{ji}) \in \mathbb{R}^{m \times n}$, where $d_{ji} \in \{0, 1\}$ takes the value 1 if the surface point s_j is covered by the dilated ball B_i and 0 otherwise.

In this paper, we define S' as the uncovered surface points, i.e.,

$$S' := \{s_j \in S \mid s_j \notin B_i, \forall p_i \in P^+\}. \quad (4)$$

Then we present the *coverage score* to each $p_i \in P^-$, which is determined by the number of samples in uncovered surface points S' that B_i covers

$$Cov_i := \sum_{j: s_j \in S'} d_{ji}. \quad (5)$$


If the dilated ball B_i of a candidate point p_i can cover lots of surface samples from the surface point set S' , p_i should be regarded as a representative candidate, which indicates that it has high (local) shape approximation capability. The coverage score favors the medial balls that dominate a larger local volume of the given shape during the selection and thus promotes a high coverage rate (See Appendix. C) of the overall shape to effectively enforce consistency with the input shape.

4.2. Uniformity Score

In addition to the coverage score that intrinsically relates the skeletal points with the model surface, one needs to guarantee uniformity among the selected skeletal points. In other words, the selected points themselves should be relatively uniformly distributed in the space instead of clustering together to avoid redundancy. For this purpose, we introduce the *uniformity score*, which aims at punishing those candidate points close to the set of currently selected points P^+ . Specifically, for each $p_i \in P^-$, its uniformity score is defined as

$$Unif_i := \text{dist}(p_i, P^+) = \min_{p_{i'} \in P^+} \|p_i - p_{i'}\|. \quad (6)$$

To illustrate the mechanics of the uniformity score defined above, consider skeletal point selection for a simple planar shape, where the p_1, p_2, p_3, p_4 are selected points within a rectangular shape. If not considering the effect of the coverage score, then according to Eq. 6, the next skeletal point p (highlighted in red) should be the intersection point of the line segments connecting p_1 and p_3 , as well as p_2 and p_4 , ensuring that $\min_{i \in \{1, 2, 3, 4\}} \|p - p_i\|$ is the



largest within the rectangle, leading to the highest possible uniformity score for the selected point. Following this pattern, the subsequent four selected points (highlighted in blue) should be placed at the centroids of their respective triangles, resulting in a relatively uniform distribution of the selected skeletal points.

4.3. Centrality Score

We further encourage the centrality of the selected points by introducing a centrality score for each candidate skeletal point. The more central a point is in terms of medial ball positioning, the higher its score. It excludes points too close to the surface to avoid their susceptibility to surface irregularities and noise. Specifically, for each $p_i \in P^-$, we define its centrality score based on its medial radius by

$$Centr_i := -\frac{1}{r_i}, \quad (7)$$

where r_i is the radius of p_i defined in Eq. 1.

4.4. Final Score

We assign a final score to each candidate $p_i \in P^-$ by combining its coverage score, uniformity score, and centrality score. We denote the three score vectors as $\mathbf{Cov} = (Cov_i)$, $\mathbf{Unif} = (Unif_i)$, and $\mathbf{Centr} = (Centr_i)$, respectively. Since these scores have different scales, we apply standardization before integrating them. Specifically, for each $p_i \in P^-$, we perform the following transformation to Cov_i , $Unif_i$ and $Centr_i$

$$\begin{aligned} Cov_i &\leftarrow \frac{Cov_i - \text{mean}(\mathbf{Cov})}{\text{std}(\mathbf{Cov})}, \\ Unif_i &\leftarrow \frac{Unif_i - \text{mean}(\mathbf{Unif})}{\text{std}(\mathbf{Unif})}, \\ Centr_i &\leftarrow \frac{Centr_i - \text{mean}(\mathbf{Centr})}{\text{std}(\mathbf{Centr})}, \end{aligned} \quad (8)$$

where $\text{mean}(\cdot)$ and $\text{std}(\cdot)$ represent the sample mean and sample standard deviation respectively. Then we integrate the three scores by summing them up, i.e.,

$$Score_i := Cov_i + Unif_i + Centr_i. \quad (9)$$

After computing $Score_i$ for all $p_i \in P^-$, we select the candidate that achieves the highest score. An overview of the heuristic algorithm is summarized in Algorithm 1. An ablation study on different scoring schemes is given in Sec. 6.3.

Algorithm 1: Coverage Axis++ for inner point selection

Input: sampled surface points S , candidate skeletal points P , dilation factor δ_r , target number of selected points $|V|$.

Result: the set of selected skeletal points P^+

Compute matrix \mathbf{D} based on Eq. 3, Initialize $S' = S$, $P^+ = \emptyset$, $P^- = P$ and $k = 1$;

while $S' \neq \emptyset$ and $k \leq |V|$ **do**

for $p_i \in P^-$ **do**

 Compute Cov_i according to Eq. 5;

 Compute $Unif_i$ according to Eq. 6;

 Compute $Centr_i$ according to Eq. 7;

end

 Compute the mean and standard deviation for the two vectors $\mathbf{Cov} = (Cov_i)$ and $\mathbf{Unif} = (Unif_i)$;

for $p_i \in P^-$ **do**

 Standardize Cov_i , $Unif_i$ and $Centr_i$ according to Eq. 8;

 Compute $Score_i$ according to Eq. 9;

end

 Select $p_{ik} \leftarrow \arg \max_{p_i \in P^-} \{Score_i\}$;

 Update $P^+ \leftarrow P^+ \cup \{p_{ik}\}$, $P^- \leftarrow P^- \setminus \{p_{ik}\}$;

 Update S' according to Eq. 4;

 Set $k \leftarrow k + 1$;

end

4.5. Candidate and Surface Point Generation

In this paper, for mesh or point cloud inputs, we always randomly generate the set of candidate points P ($|P| = 10000$) inside the volume using volumetric sampling[†]. The number of surface samples S is 1500 in all the comparisons (except for Appendix. A). For mesh input, we employ Poisson-disk sampling [CCS12] to get the surface sample. For point cloud input, surface samples are obtained by white noise sampling [Jac*21], and fast winding number [BDS*18] is used for inside-outside determination.

4.6. Connectivity of Skeletal Points

4.6.1. Mesh Input

For a mesh input, we aim to build up a connection while preserving the topology of the skeleton, for which we employ the edge collapse strategy of Q-MAT [LWS*15] which is based on a Voronoi diagram. Following [ACK01; WWX*22; DLX*22], we initialize the initial Voronoi diagram \mathbf{VD} using the surface points. We then embed the selected points in the Voronoi diagram using the nearest projection, following [DLX*22]. We simplify \mathbf{VD} using the edge collapse strategy [LWS*15]. Specifically, during this process, each edge is checked and collapsed using the edge collapse strategy [LWS*15] based on the Spherical Quadric Error Metric [LWS*15; TGB13]. We use additional conditions during edge collapse: If both endpoints of an edge to be collapsed are not





Figure 3: An illustration of the edge collapse process. The surface is in black, with selected skeletal points highlighted in red. The collapsed edges are shown with dashed lines.

selected points, we directly apply the edge collapse strategy following [LWS*15]. If one endpoint of an edge is a selected point and the other is not, during edge collapse, the selected point is preserved. If both endpoints of an edge are selected points, the edge is preserved and skipped. This process continues until all edges have endpoints that belong to the selected points; See Fig. 3. We adopt topology preservation [DEGN99] and mesh inversion avoidance strategy [GH97] during the collapse. The simplification result \mathbf{VD}' serves as the final MAT of the given mesh input.

Remark. Notably, not all Voronoi vertices inside the given shape can be used to approximate the medial surface. One has to use *Poles* (a subset of inside Voronoi vertices) for medial surface approximation according to [ACK01]. However, in practice, we find the collapse over the original Voronoi diagram \mathbf{VD} yields relatively stable outcomes for a compact medial surface approximation. Furthermore, the connectivity of skeletal points (edges and facets in MAT) is crucial information in the output MAT results. This connectivity is significantly important for downstream applications such as shape analysis and reconstruction, which we discuss in detail in Appendix E.

4.6.2. Point Cloud Input

Topology preservation during connection establishment for point clouds is challenging. The aforementioned edge collapse techniques [LWS*15; PWG*19] are not applicable as they require a watertight surface. Thus, we approximate the MAT connection using the Power Diagram following [ACK01]. We compute the Power diagram \mathbf{PD} on the set of $\{B^l(p_i, r'_i), p_i \in P^+\} \cup \{B(s_j, \delta_r), s_j \in S\}$ where P^+ , S and $r'_i = r_i + \delta_r$ (Eq. 2) are sets of selected points, surface points and dilated radius, respectively. Then, the connectivity among skeletal points is obtained by extracting edges among the selected inner points on the dual of \mathbf{PD} , regular triangulation \mathbf{RT} (in other words, weighted alpha complexes). In the inset figure, the P^+ points are depicted in red, while the surface samples are shown in black. The medial balls with original radius r and dilated radius r' are illustrated in red and yellow, respectively. Connectivity is established by the connections among selected points of the regular triangulation \mathbf{RT} , which are also colored as black solid lines. Additionally, the power plane between the two sites is visualized in blue. This approach can also be applied to mesh inputs, as the computation of the Power Diagram does not depend on the underlying surface being watertight. An overview of our pipeline is shown in Fig. 2.



[†] <https://www.meshlab.net>

4.7. Discussion

4.7.1. Complexity

Our method (Algorithm 1) has a polynomial complexity. Specifically, computing Cov_i and $Centr_i$ costs $O(|S|)$ operations, while $Unif_i$ requires $O(|V|)$ operations. Hence summing up for all $p_i \in P^-$ gives complexity $O(|P| \cdot |S|)$, considering that $|V|$ is usually very small compared to $|S|$ and $|P|$. Standardizing **Cov**, **Unif** and **Centr**, as well as sorting, each take $O(|P|)$ operations. Taking into account that the procedure repeats for a maximum of $|V|$ times, the overall complexity of Algorithm 1 can be expressed as $O(|S| \cdot |P| \cdot |V|)$. Fig. 12 validates this complexity empirically by illustrating the impact of different set sizes $|P|$, $|S|$ and $|V|$ on the running time of Coverage Axis++. In contrast, Coverage Axis requires an exponential cost in the worst case. Although modern solvers accelerate the solving process, long computation time is incurred in the worst case; See Tab. 1.

4.7.2. More Features

Coverage Axis++ offers control over the number of selected inner points compared to Coverage Axis. In Sec. 6.1, we explore its robustness to various target numbers of selected skeletal points. This ability to specify a specific number of skeletal points is crucial for learning-based methods, ensuring compatibility with algorithms that assume a fixed number of input points [HWQ*19; THT*21; HCY*22; PGT*24]. Unlike Coverage Axis [DLX*22] or Neural Skeleton [CD23], which enforces a hard constraint requiring all surface samples to be covered, our method relaxes this constraint, enhancing robustness to coverage conditions. This relaxation is particularly advantageous when dealing with a large number of surface samples. Coverage Axis++ maintains a high coverage rate of the original shape, as demonstrated in Appendix. C.

Coverage Axis++ does not impose strict requirements on the watertightness or manifoldness of the input compared with methods like Q-MAT [LWS*15; PWG*19], making it versatile for handling meshes and point clouds.

Additionally, our method only requires two hyperparameters: target skeletal number $|V|$, and dilation factor δ_r .

5. Experimental Results

We conduct extensive qualitative and quantitative evaluations of the proposed method to demonstrate its effectiveness comprehensively. All experiments are conducted on a computer with a 4.00 GHz Intel(R) Core(TM) i7-6700K CPU and 32 GB memory. The sizes of all the models are normalized to the $[0, 1]$ range.

Given the target skeletal point number, Coverage Axis++ only includes one hyperparameter: δ_r (Eq. 2). All the experimental results are generated using a consistent parameter setting, i.e., for all kinds of input, we always set $\delta_r = 0.02$. In experiments where we compare our method with Coverage Axis, we set the number of selected points $|V|$ for Coverage Axis++ to be equal to the number obtained from the SCP solution for the Coverage Axis. We provide additional analyses on parameters δ_r and $|V|$ in Sec. 6.

As suggested by [DLX*22], we employ the two-sided Hausdorff

distance (HD) to assess the consistency between the original and reconstructed surfaces. We use \overrightarrow{e} for surface-to-reconstruction HD, \overleftarrow{e} for reconstruction-to-surface HD, and \overleftrightarrow{e} for two-sided HD. The reconstructed surfaces are obtained by interpolating medial balls based on skeleton connectivity, akin to [LWS*15; DLX*22; WWX*22]. Errors are normalized by the diagonal length of the input model bounding box. In quantitative results, the best results for runtime and accuracy are indicated with blue and orange highlighting, respectively. A gallery of skeletonization results of mesh and point cloud inputs are shown in Fig. 1 and Fig. 5, respectively.

5.1. Mesh Inputs

5.1.1. Comparison with Coverage Axis

As Coverage Axis [DLX*22] can handle both mesh and point cloud inputs, we compared the proposed method with it on both two inputs. The quantitative results are summarized in Tab. 1. As can be observed, for both mesh and point cloud inputs, Coverage Axis++ requires less runtime while maintaining competitive accuracy on most shapes. We observed that in specific cases, such as the Vase model, the Coverage Axis shows a speed advantage over our approach. This advantage is primarily attributed to the efficiency of its heuristic algorithm[‡] for solving the Set Cover Problem (SCP). However, such instances are infrequent due to the theoretical computational complexity associated with SCP. When evaluating the overall performance, our method significantly outperforms the Coverage Axis in terms of running time, where ours achieves 5.362% reconstruction error in 2.9s on average, whereas the Coverage Axis achieves 6.574% reconstruction error in more than 236.1s. The same advantage is also reflected in point cloud inputs. Actually, the runtime for Coverage Axis tends to be longer for shapes containing planar structures due to the complexity of the algorithm; See discussion in Appendix D. The running times of the Coverage Axis are even longer when using more surface samples; See Appendix. A. A qualitative comparison is provided in Fig. 6.

5.1.2. Comparison with MATFP

We further compare our method with the recent work MATFP [WWWG22]. MATFP aims to compute MAT with feature preservation, and it only takes watertight mesh input. Besides, similar to the Coverage Axis, the user cannot explicitly specify the final number of skeletal points. As shown in Fig. 4, we find MATFP typically yields results with redundant points. In comparison, Coverage Axis++ produces more compact results while using less computation time (13.11s v.s. ours 2.61s). Since MATFP is a multi-stage approach, we detail the running time of each stage in Appendix B.

5.1.3. Comparison with Q-MAT and Q-MAT+

We compare our method with two representative methods for MAT computation given the mesh inputs: Q-MAT [LWS*15] and its variant Q-MAT+ [PWG*19]. Both two methods take watertight meshes as inputs and cannot handle point clouds or polygon soups. Tab. 2

[‡] Mixed-integer linear programming in SciPy or MATLAB typically employs heuristics to search for integer-feasible solutions.

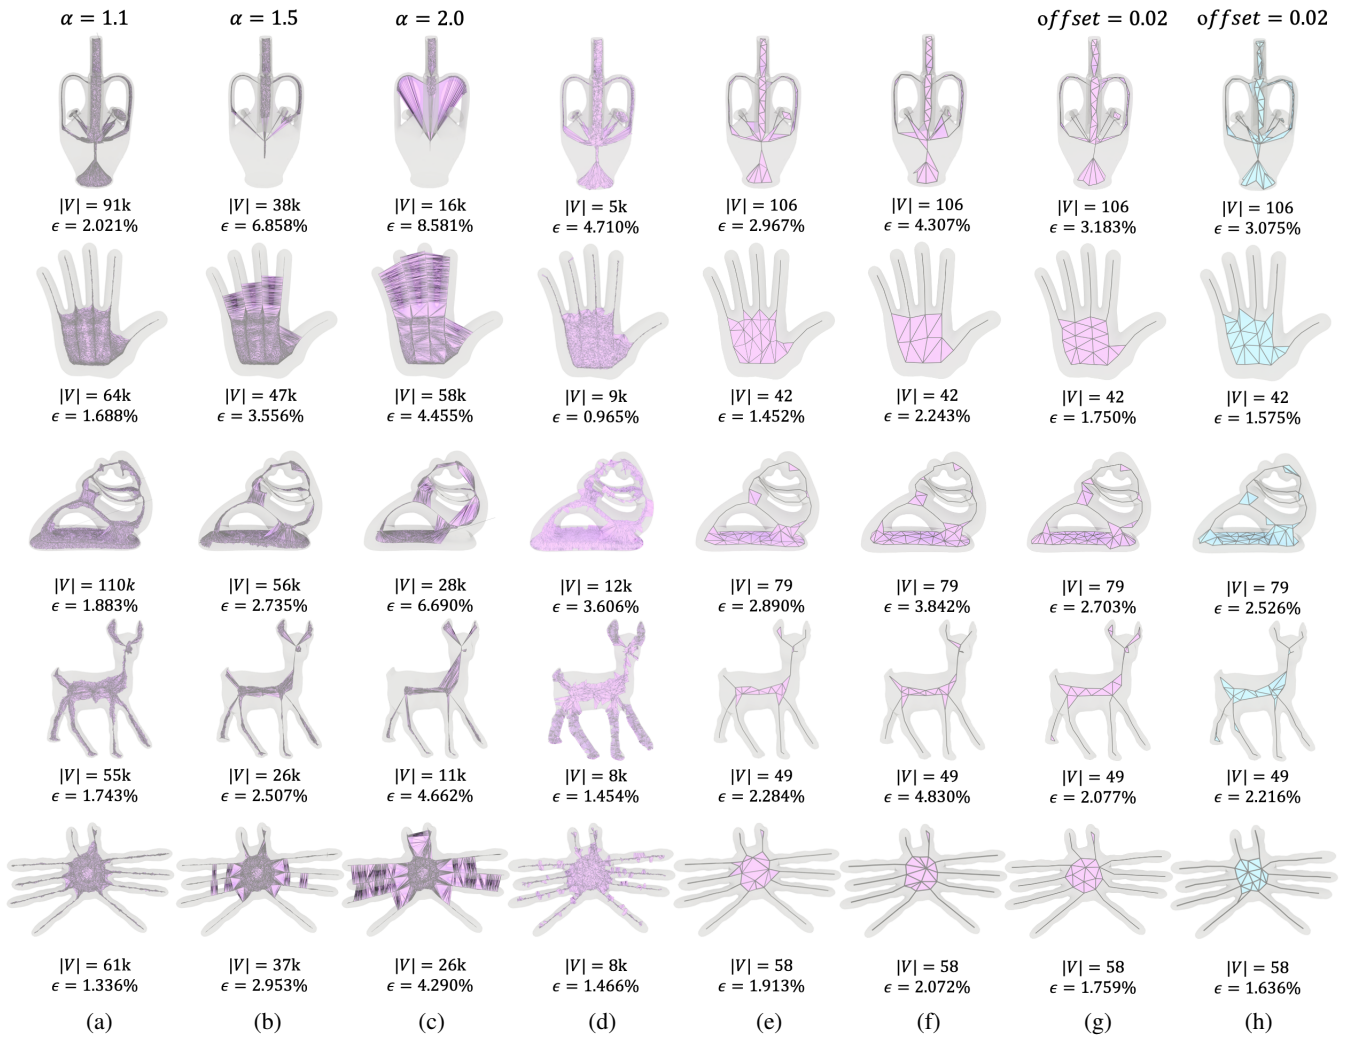


Figure 4: Qualitative comparisons. (a-c) SAT. (d) MATFP. (e) QMAT. (f) QMAT+. (g) Coverage Axis. (h) Coverage Axis++.




Figure 5: A gallery of skeletonization results of point clouds using Coverage Axis++.




Figure 6: Comparison with existing shape skeletonization methods for point clouds.

shows that Coverage Axis++ performs comparably or superior to Q-MAT and Q-MAT+ in terms of reconstruction accuracy and running time. Note that both Q-MAT and Q-MAT+ require a watertight mesh for Voronoi initialization, while our method supports *randomly* generated candidate points inside the volume.

5.1.4. Comparison with SAT

SAT [MGP10] encounters a major drawback as it tends to favor a dense representation with numerous vertices, hindering the generation of a simple and compact skeleton of the given shape. Note SAT only excels specifically with watertight surface meshes. We test with different scaling factors, $\alpha = 1.1, 1.5,$ and 2.0 , revealing that with a smaller α , SAT achieves high accuracy but results in a representation with a large number of vertices (Fig.4(a-c)). Us-

ing a larger α for higher abstraction compromises shape structure and introduces significant approximation errors. As a comparison, Coverage Axis++ achieves a balance between fidelity, efficiency, and compression abilities.

5.2. Point Cloud Inputs

5.2.1. Comparison with Neural Skeleton

We compare Coverage Axis++ with Neural Skeleton [CD23] given the point cloud inputs. Neural Skeleton first trains a SDF w.r.t. the input point cloud, then samples uniform surface points and skeletal points from the implicit field. Finally, the skeleton mesh is recovered using a cover-set formulation [DLX*22] solved as a Mixed-Integer Linear Program. The implicit field improves the robustness of the Neural Skeleton in handling challenging inputs. As

Table 3: Comparison on shape approximation errors among Point2Skeleton (P2S), Deep Point Consolidation (DPC), Neural Skeleton, Coverage Axis and Coverage Axis++.

Model	P2S $\epsilon \leftrightarrow$		DPC $\epsilon \leftrightarrow$		Neural Skeleton $\epsilon \leftrightarrow$		Coverage Axis $\epsilon \leftrightarrow$		Coverage Axis++ $\epsilon \leftrightarrow$	
	$ V $	ϵ	$ V $	ϵ	$ V $	ϵ	$ V $	ϵ	$ V $	ϵ
Ant-2	100	16.412%	1194	8.863%	54	3.993%	58	2.350%	58	2.113%
Bottle	100	2.955%	1194	2.752%	13	2.390%	14	2.956%	14	2.599%
Chair-2	100	6.552%	1194	4.807%	134	3.651%	89	2.890%	89	3.012%
Dog	100	6.047%	1194	5.237%	42	3.074%	49	2.174%	49	2.306%
Dolphin	100	5.925%	1194	7.916%	36	3.269%	49	1.971%	49	1.876%
Fertility	100	8.162%	1194	4.226%	76	5.909%	79	3.428%	79	3.102%
Guitar	100	2.052%	1194	3.231%	65	2.872%	60	2.032%	60	2.344%
Hand	100	9.672%	1194	4.110%	39	4.819%	44	3.441%	44	2.328%
Kitten	100	8.724%	1194	5.332%	41	3.883%	50	3.450%	50	2.962%
Snake	100	15.021%	1194	1.736%	36	1.419%	40	1.309%	40	1.402%
Average	-	8.003%	-	4.815%	-	3.528%	-	2.515%	-	2.404%

shown in Tab. 3, our method produces results with higher accuracy. Moreover, the average running time of Coverage Axis++ and NeuralSkeleton is 1.63s v.s. 59.27s. The detailed runtime statistics of Neural Skeleton are given in Appendix B. Note that the Neural Skeleton does not support specifications for the number of skeletal points.

5.2.2. Comparison with Point2Skeleton

Compared with Point2Skeleton [LLL*21], a learning-based skeletonization method for point clouds, our approach exhibits better performance. As shown in Fig. 6 and Tab. 3, Point2Skeleton, trained on a large dataset ShapeNet [CFG*15], falls short on shapes unseen during training and struggles with higher topology complexity. However, our method demonstrates robustness, generating high-quality skeletal representations with accurate geometries and faithful structures across diverse shapes.

5.2.3. Comparison with Deep Point Consolidation

As shown in Fig. 6 and Tab. 3, Deep Point Consolidation (DPC) [WHG*15] can only produce unstructured inner points without connections and the reconstruction results exhibit large errors. Even worse, the reconstructed topologies by DPC are usually inconsistent with the original input. In contrast, our reconstruction results reach better approximation accuracy with respect to the original geometry using fewer skeletal points.

6. More Discussions

6.1. Skeletal Point Number

In Sec. 4.7.2, we highlight Coverage Axis++'s practical benefit of flexible specification for the target number of skeletal points. Fig. 7 presents our experiments with different target skeletal numbers using three models and three choices of $|V|$. Larger values of $|V|$ result in more accurate skeletal representations, as observed in the experiments where reconstruction error consistently decreases with increasing $|V|$. However, choosing excessively large $|V|$ contradicts our objective of generating a simple and compact skeleton for shape abstraction. Our results show that specifying a smaller number of skeletal points (e.g., 30 or 50) achieves simpler shape abstraction with only a slight increase in reconstruction error. The flexibility allows users to tailor the method for either a simple or more accurate model, depending on the requirement for shape abstraction.

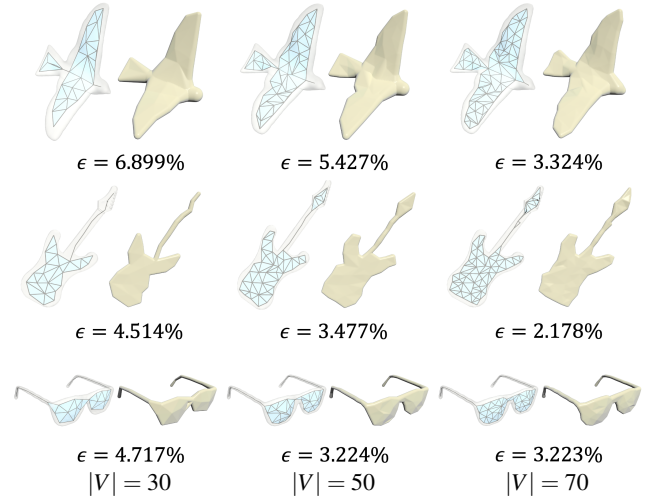


Figure 7: Influence of target numbers $|V|$. We test our model with the target numbers being 30, 50 and 70.




Figure 8: Skeletonization results using different dilation offsets.

6.2. Dilation Factor

Next, we perform an ablation study on the dilation factor δ_r using the Femur model as an example. As shown in Fig. 8, increasing δ_r allows balls to cover more surface samples, simplifying the skeletal representation. We always set $|V| = 30$ in Algorithm 1, but when δ_r is 0.05 or larger, the algorithm terminates with fewer than 30 selected points as dilated balls already cover all surface samples. A larger δ_r simplifies the representation but sacrifices geometric details, compromising shape structure preservation. Conversely, an excessively small δ_r contradicts simplicity and compactness. We maintain consistency by setting δ_r to 0.02 in all experiments, balancing geometric details and skeletal compactness.

6.3. Ablation on Scoring Terms

We evaluate our method with different scoring terms defined as different combination schemes of the *coverage* (Eq. 5), *uniformity* (Eq. 6) and *centrality* (Eq. 7) scores on the Crab and Human models, where we set $|V| = 60$ for both models. From Fig. 10, we observe that the inclusion of the uniformity score results in skeletonizations with fewer sharp triangles, indicating a more uniform distribution of selected points. Furthermore, incorporating the centrality score enables the selected skeletal points to capture more accurate MAT information, leading to a more compact and regular skeletonization. Consequently, incorporating all three scores, as shown in Eq. 9

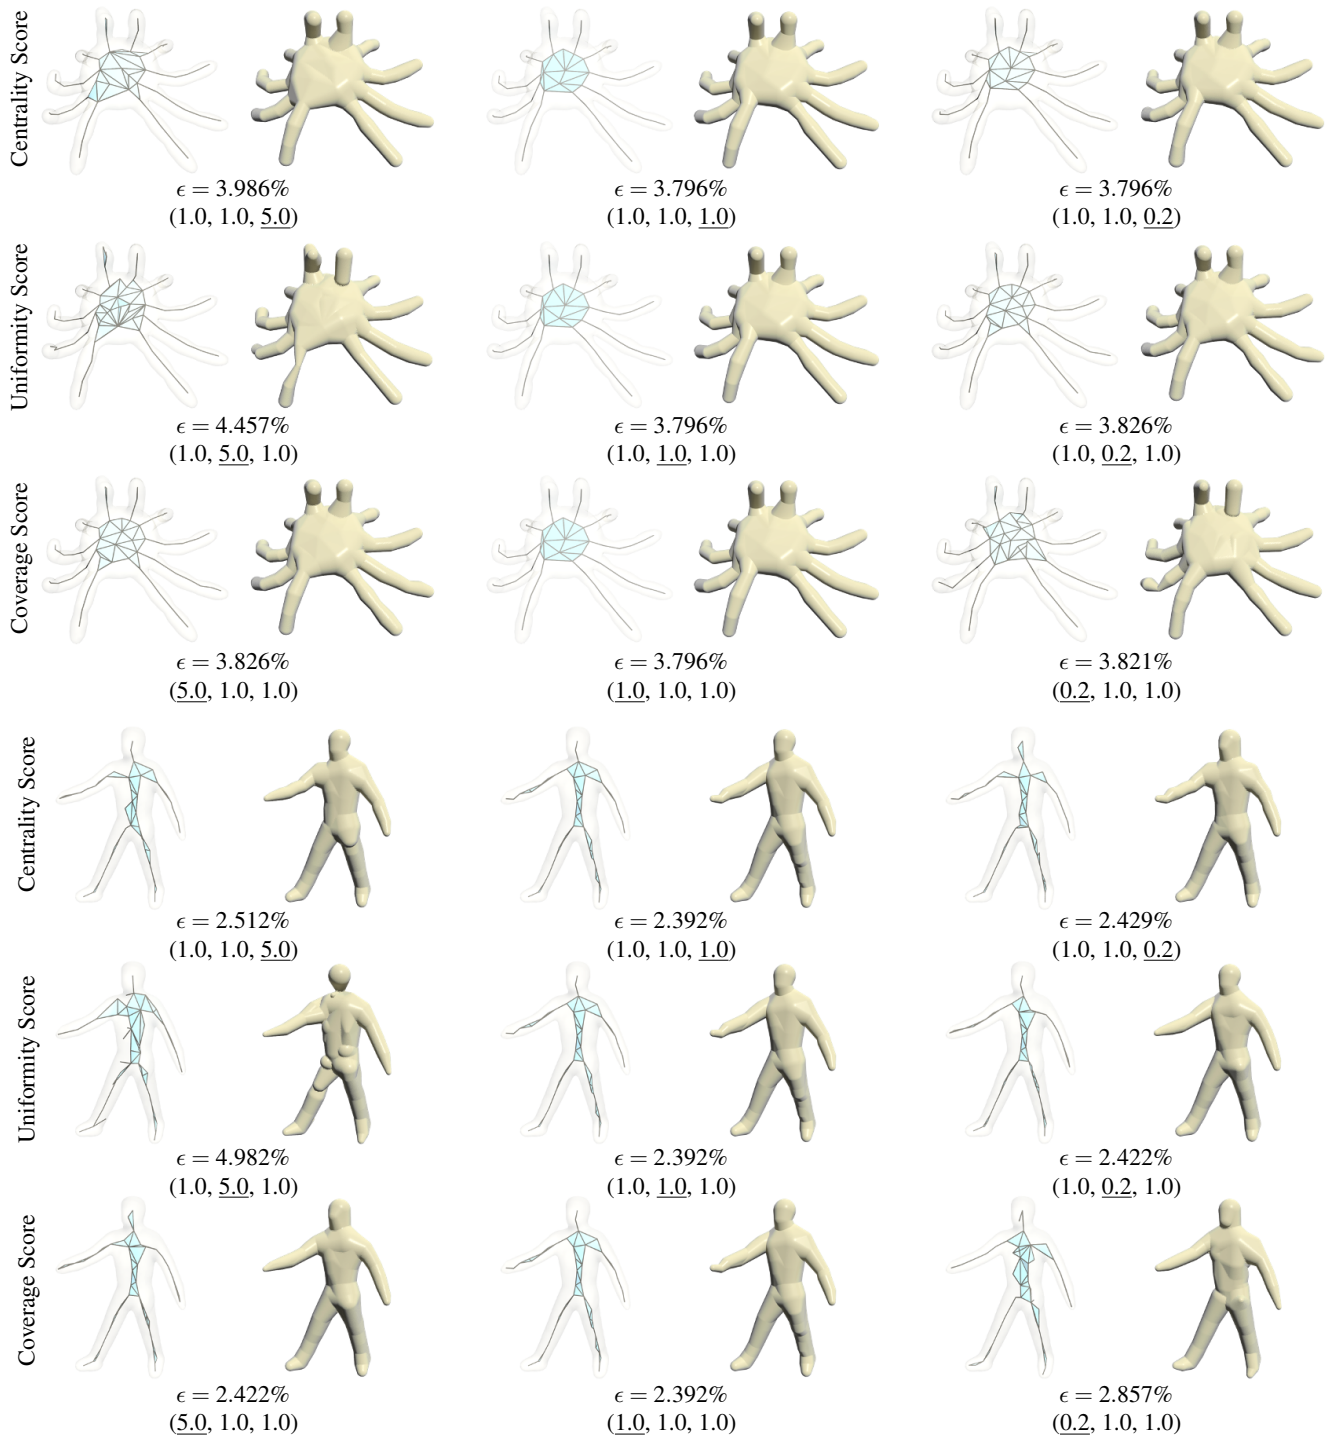


Figure 9: We conducted ablation experiments on different weights. We test our method using different weights corresponding to the centrality score, uniformity score, and coverage score, e.g., (1.0, 1.0, 0.2) refers to $Score_i = 1.0 \cdot Cov_i + 1.0 \cdot Uni_f_i + 0.2 \cdot Centr_i$. We visualize the skeletal results and evaluate the reconstruction error. We test with Human and Crab models.

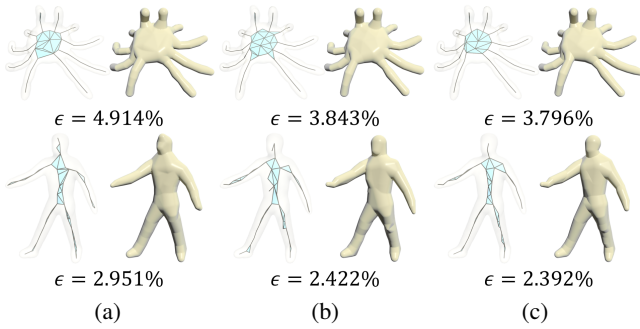


Figure 10: Performance of different combinations of the three scores. (a) coverage score. (b) coverage + uniformity score. (c) coverage + uniformity + centrality score (Eq. 9).

yields improved skeletonization outcomes, i.e., compact and high-accuracy shape approximation.

Although using equal weights for the terms in Eq.9 effectively achieves compact skeletonization and high reconstruction accuracy in our main experiments without requiring special weight adjustments, we investigate the impact of different weighting schemes for each term. Concretely, we scale the weight of each term down and up by a factor of 5, while keeping the weights of the other terms unchanged, to investigate the influence of different terms. We test our algorithm using two models: Crab and Human. As shown in Fig. 9, we observe that applying a moderate uniformity weight significantly enhances the spatial distribution uniformity of the model vertices. However, using excessively large centrality score weights may deteriorate the regular structure of skeletal points, resulting in increased reconstruction error.

6.4. Shape Skeletonization for Triangle Soups

We evaluate our method for shape skeletonization using triangle soup inputs. For the experiment, we apply a 0.25% perturbation relative to the bounding box to each vertex of the triangles. The skeletonization process is identical to that used for handling point cloud inputs. In this experiment, we sample 1500 surface points from the soup and use 10000 randomly generated candidate points inside the volume, as we described in Sec. 4.5. As shown in Fig. 11, our method yields high-quality skeletonization and shape approximation results with the triangle soup inputs.

6.5. Running Time Analysis

The overall complexity of Algorithm 1 is $O(|S| \cdot |P| \cdot |V|)$. Next, we validate the practical running time complexity of our method using different set sizes $|P|$, $|S|$, and $|V|$. As shown in Fig. 12, Coverage Axis++ has polynomial complexity relationships among various variables, maintaining overall low computational costs.


6.6. Limitations and Future Works


Coverage Axis++ approximates the whole shape by an abstraction of local geometry with the union of the local geometries forming the entire shape. One limitation of Coverage Axis++ is that its whole process is conducted without a




Figure 11: Shape skeletonization results for triangle soup. (a) Input triangle soup. (b) Skeletonization result. (c) Shape reconstruction.

hard constraint forcing the consistency between the selected medial balls and the shape surface (See coverage rate statistics in Appendix. C), like that in Coverage Axis [DLX*22]. Although local approximation for shape skeletonization has been evidenced by some previous works [LLL*21; GYY*23; ACK01], we acknowledge that the inclusion of a global constraint imposes a much stricter guarantee for the shape correspondence. Since the goal of this paper is to compute an overall compact skeletonization result, our method may encounter challenges when handling shapes like CAD models that have extremely thin structures or sharp on-surface features (see the inset figure). In the future, introducing additional terms to preserve sharp features during shape skeletonization would be promising. For example, an algorithm can be developed to build up the connection of sharp features and the medial axis, ensuring the skeletonization result retains the sharp features on the surface. Our method can generate a compact representation in terms of shape coverage using medial balls as primitives. We admit that while minimizing the redundancy of skeletal points is highly desirable, there is still potential to further reduce the number of skeletal points. For example, planar components such as the body of a guitar, the seat of a chair, and the base of a fertility statue are currently oversam-






(a) Running time w.r.t. target skeletal point numbers $|V|$ with $|S| = 2000$, $|P| = 10000$.



(b) Running time w.r.t. surface sample numbers $|S|$ with $|V| = 100$, $|P| = 10000$.



(c) Running time w.r.t. candidate skeletal point numbers $|P|$ with $|V| = 100$, $|S| = 2000$.

Figure 12: Runtime performance w.r.t. different parameter settings.

pled because the coverage capability is limited by the medial balls that are isotropic, as shown in Fig. 1. Future work could explore the use of different types of bounding volume primitives, such as slab meshes [LWS*15], ellipsoids [LCWK07], or medial skeletal diagrams [GWM23] instead of medial balls, to achieve a more compact representation with fewer skeletal points, while maintaining the coverage optimization goal.

Next, we discuss more potential applications and future works. In this paper, we present an effective and efficient method for skeletal point selection and MAT computation. The crucial geometrical and structural information encoded in the MAT could be helpful in various downstream applications. For instance, in the task of shape segmentation [LLL*20], the MAT can be leveraged for identifying various types of junctions between different parts of a 3D shape using information like manifold branch and dimensional change of the MAT without relying on 3D ground truth supervision for segmentation. In the future, we believe that a part-level shape prior derived from the MAT through self-supervised learning could benefit part-aware shape generation and reconstruction [LLL*24]. Meanwhile, the discriminative feature represented by the MAT could also facilitate the 3D shape recognition as shown by [HWQ*19]. MAT has demonstrated its effectiveness in deep, topology-aware generation for 3D shapes [PGT*24] with the neural skeleton-based representation encoding information on both shape topology and geometry. It has also been evidenced that the topological and geometrical information of the medial mesh can be useful in remeshing and vol-

umetric subdivision, e.g., achieving all-hexahedral mesh generation of the given shape, making the method more general than previous methods based on curve-skeleton [ZXW*24].

Even the point selection part of our method could serve as an effective point sampling strategy for compact shape encoding in wide applications. For example, the point selection method [DLX*22] has been shown to be effective in enhancing computational and memory efficiency for 3D shape representation, e.g., TetSphere Splatting [GWHM24]. More potential applications of our point selection method can be those point-based methods. Specifically, one may achieve compact 2D/3D Gaussian Splatting [HYC*24; KKLD23], also known as Gaussian pruning, given a properly defined coverage score during the training. Actually, our point selection algorithm can serve as a general sampling approach applicable to various tasks. As long as an original input representation can be expressed as *candidate points* and an associated *covering matrix* are properly defined, thanks to the SCP, our method could yield a more *compact representation* for the given *redundant input*.

7. Conclusion

In this paper, we present Coverage Axis++, an efficient shape skeletonization method. Unlike existing methods that depend on input watertightness or suffer from high computational costs, Coverage Axis++ addresses these challenges by employing a heuristic algorithm considering shape coverage, uniformity, and centrality for selecting skeletal points, which ensures an efficient approximation of the MAT while significantly saving the computational cost. Compared with existing methods, our simple yet effective approach produces a compact representation of MAT, offers versatility in skeletonizing various shape representations, allows customizable specification of skeletal points, and achieves highly efficient computation with improved reconstruction accuracy. Extensive experiments across diverse 3D shapes demonstrate the efficiency and effectiveness of Coverage Axis++.

8. Acknowledgements

The authors would like to thank Baorong Yang, Mengyuan Ge, and Ningna Wang for the fruitful discussion and the anonymous reviewers for their valuable comments and suggestions. Zhiyang Dou is supported by the Innovation and Technology Commission of the HKSAR Government under the InnoHK initiative. Shiqing Xin is supported by the National Natural Science Foundation of China (Ref: 62272277). This research is partly supported by Innovation and Technology Commission (Ref: ITS/319/21FP) and Research Grant Council (Ref: 17210222), Hong Kong.

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, 12.

- [Ale21] ALEXA, MARC. “Polycover: Shape approximating with discrete surface orientation”. *IEEE Computer Graphics and Applications* 41.3 (2021), 85–95 2.
- [And09] ANDREA TAGLIASACCHI Hao Zhang, DANIEL COHEN-OR. “Curve skeleton extraction from incomplete point cloud”. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 28.3 (2009), 71:1–71:9 3.
- [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 2.
- [Aur87] AURENHAMMER, FRANZ. “Power diagrams: properties, algorithms and applications”. *SIAM Journal on Computing* 16.1 (1987), 78–96 3.
- [BDS*18] BARILL, GAVIN, DICKSON, NEIL G, SCHMIDT, RYAN, et al. “Fast winding numbers for soups and clouds”. *ACM Transactions on Graphics (TOG)* 37.4 (2018), 1–12 5.
- [Blu*67] BLUM, HARRY et al. *A transformation for extracting new descriptors of shape*. Vol. 43. MIT press Cambridge, MA, 1967 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.
- [CCS12] CORSINI, MASSIMILIANO, CIGNONI, PAOLO, and SCOPIGNO, ROBERTO. “Efficient and flexible sampling with blue noise properties of triangular meshes”. *IEEE transactions on visualization and computer graphics* 18.6 (2012), 914–924 5.
- [CD23] CLÉMOT, MATTÉO and DIGNE, JULIE. “Neural skeleton: Implicit neural representation away from the surface”. *Computers & Graphics* 114 (2023), 368–378 1–3, 6, 9.
- [CFG*15] CHANG, ANGEL X, FUNKHOUSER, THOMAS, GUIBAS, LEONIDAS, et al. “Shapenet: An information-rich 3d model repository”. *arXiv preprint arXiv:1512.03012* (2015) 10.
- [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.
- [DEGN99] DEY, TAMAL, EDELSBRUNNER, HERBERT, GUHA, SUMANTA, and NEKHAYEV, DMITRY. “Topology preserving edge contraction”. *Publications de l’Institut Mathématique* 66 (1999) 5.
- [DLX*22] DOU, ZHIYANG, LIN, CHENG, XU, RUI, et al. “Coverage axis: Inner point selection for 3d shape skeletonization”. *Computer Graphics Forum*. Vol. 41. 2. Wiley Online Library, 2022, 419–432 1–3, 5–7, 9, 12, 13.
- [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 2.
- [DWL*23] DOU, ZHIYANG, WU, QINGXUAN, LIN, CHENG, et al. “Tore: Token reduction for efficient human mesh recovery with transformer”. *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2023, 15143–15155 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* 27.10 (2020), 3982–3993 1, 2.
- [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.
- [FXX*22] FU, ZHIYING, XU, RUI, XIN, SHIQING, et al. “Easyvrmodelling: Easily create 3d models by an immersive vr system”. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 5.1 (2022), 1–14 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 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.
- [GWHM24] GUO, MINGHAO, WANG, BOHAN, HE, KAIMING, and MATUSIK, WOJCIECH. *TetSphere Splatting: Representing High-Quality Geometry with Lagrangian Volumetric Meshes*. 2024. arXiv: 2405.20283 [cs.CV] 2, 13.
- [GWM23] GUO, MINGHAO, WANG, BOHAN, and MATUSIK, WOJCIECH. “Medial Skeletal Diagram: A Generalized Medial Axis Approach for Compact 3D Shape Representation”. *arXiv preprint arXiv:2310.09395* (2023) 2, 13.
- [GYY*23] GE, MENGYUAN, YAO, JUNFENG, YANG, BAORONG, et al. “Point2MM: Learning medial mesh from point clouds”. *Computers & Graphics* 115 (2023), 511–521 2, 3, 12.
- [HCY*22] HU, JIANWEI, CHEN, GANG, YANG, BAORONG, et al. “IMMAT: Mesh reconstruction from single view images by medial axis transform prediction”. *Computer-Aided Design* 150 (2022), 103304 2, 6.
- [HWC*13] HUANG, HUI, WU, SHIHAI, COHEN-OR, DANIEL, et al. “L1-medial skeleton of point cloud.” *ACM Trans. Graph.* 32.4 (2013), 65–1 3.
- [HWQ*19] HU, JIANWEI, WANG, BIN, QIAN, LIHUI, et al. “MAT-Net: Medial Axis Transform Network for 3D Object Recognition.” *IJCAI*. 2019, 774–781 2, 6, 13.
- [HWY*23] HU, JIANWEI, WANG, NINGNA, YANG, BAORONG, et al. “S3DS: Self-supervised Learning of 3D Skeletons from Single View Images”. *Proceedings of the 31st ACM International Conference on Multimedia*. 2023, 6948–6958 2.
- [HYC*24] HUANG, BINBIN, YU, ZEHAI, CHEN, ANPEI, et al. “2d gaussian splatting for geometrically accurate radiance fields”. *arXiv preprint arXiv:2403.17888* (2024) 13.
- [Jac*21] JACOBSON, ALEC et al. *gptoolbox: Geometry Processing Toolbox*. <http://github.com/alecjacobson/gptoolbox>. 2021 5.
- [KKLD23] KERBL, BERNHARD, KOPANAS, GEORGIOS, LEIMKÜHLER, THOMAS, and DRETTAKIS, GEORGE. “3d gaussian splatting for real-time radiance field rendering”. *ACM Transactions on Graphics* 42.4 (2023), 1–14 13.
- [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.
- [LCWK07] LU, LIN, CHOI, YI-KING, WANG, WENPING, and KIM, MYUNG-SOO. “Variational 3D shape segmentation for bounding volume computation”. *Computer Graphics Forum*. Vol. 26. 3. Wiley Online Library, 2007, 329–338 13.
- [LFWN20] LIN, CHENG, FAN, TINGXIANG, WANG, WENPING, and NIESSNER, MATTHIAS. “Modeling 3d shapes by reinforcement learning”. *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X 16*. Springer, 2020, 545–561 2.
- [LGC*23] LIU, HSUEH-TI DEREK, GILLESPIE, MARK, CHISLETT, BENJAMIN, et al. “Surface simplification using intrinsic error metrics”. *ACM Transactions on Graphics (TOG)* 42.4 (2023), 1–17 2.
- [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 3.

- [LL*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, 2, 13.
- [LL*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 1, 3, 10, 12.
- [LL*24] LIU, ANRAN, LIN, CHENG, LIU, YUAN, et al. “Part123: Part-aware 3D Reconstruction from a Single-view Image”. *arXiv preprint arXiv:2405.16888* (2024) 2, 13.
- [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 1–3, 5, 6, 13.
- [LXC*21] LI, JIEFENG, XU, CHAO, CHEN, ZHICUN, et al. “Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation”. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2021, 3383–3393 1.
- [LYZ*23] LIN, GUYING, YANG, LEI, ZHANG, CONGYI, et al. “Patch-Grid: An Efficient and Feature-Preserving Neural Implicit Surface Representation”. *arXiv preprint arXiv:2308.13934* (2023) 2.
- [LZG*23] LI, CHUNHUI, ZHOU, MINGQUAN, GENG, GUOHUA, et al. “EPCS: Endpoint-based part-aware curve skeleton extraction for low-quality point clouds”. *Computers & Graphics* 117 (2023), 209–221 3.
- [MGP10] MIKLOS, BALINT, GIESEN, JOACHIM, and PAULY, MARK. “Discrete scale axis representations for 3D geometry”. *ACM SIGGRAPH 2010 papers*. 2010, 1–10 2, 9.
- [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.
- [MSLJ23] MARSCHNER, ZOË, SELLÁN, SILVIA, LIU, HSUEH-TI DEREK, and JACOBSON, ALEC. “Constructive Solid Geometry on Neural Signed Distance Fields”. *SIGGRAPH Asia 2023 Conference Papers*. 2023, 1–12 2.
- [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 2.
- [PGT*24] PETROV, DMITRY, GOYAL, PRADYUMN, THAMIZHARASAN, VIKAS, et al. “GEM3D: Generative Medial Abstractions for 3D Shape Synthesis”. *arXiv preprint arXiv:2402.16994* (2024) 2, 6, 13.
- [PWG*19] PAN, YILING, WANG, BIN, GUO, XIAOHU, et al. “Q-MAT+: An error-controllable and feature-sensitive simplification algorithm for medial axis transform”. *Computer Aided Geometric Design* 71 (2019), 16–29 1–3, 5, 6.
- [RAV*19] REBAIN, DANIEL, ANGLAS, 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) 2.
- [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.
- [SKSJ20] SELLÁN, SILVIA, KESTEN, JACOB, SHENG, ANG YAN, and JACOBSON, ALEC. “Opening and closing surfaces”. *ACM Transactions on Graphics (TOG)* 39.6 (2020), 1–13 2.
- [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 2, 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 2.
- [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 5.
- [THP*19] TANG, JIAPENG, HAN, XIAOGUANG, 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.
- [THT*21] TANG, JIAPENG, HAN, XIAOGUANG, TAN, MINGKUI, et al. “Skeletonnet: A topology-preserving solution for learning mesh reconstruction of object surfaces from rgb images”. *IEEE transactions on pattern analysis and machine intelligence* 44.10 (2021), 6454–6471 2, 6.
- [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, 10.
- [WHS*24] WANG, NINGNA, HUANG, HUI, SONG, SHIBO, et al. “MAT-Topo: Topology-preserving Medial Axis Transform with Restricted Power Diagram”. *arXiv preprint arXiv:2403.18761* (2024) 3.
- [WTJL24] WEN, QINGMENG, TAFRISHI, SEYED AMIR, JI, ZE, and LAI, YU-KUN. “GLSkeleton: A geometric Laplacian-based skeletonisation framework for object point clouds”. *IEEE Robotics and Automation Letters* (2024) 2.
- [WWWG22] WANG, NINGNA, WANG, BIN, WANG, WENPING, and GUO, XIAOHU. “Computing medial axis transform with feature preservation via restricted power diagram”. *ACM Transactions on Graphics (TOG)* 41.6 (2022), 1–18 2, 3, 6.
- [WWX*22] WANG, PENGFEI, WANG, ZIXIONG, XIN, SHIQING, et al. “Restricted delaunay triangulation for explicit surface reconstruction”. *ACM Transactions on Graphics* 41.5 (2022), 1–20 5, 6.
- [WYT23] WEN, CHENG, YU, BAOSHENG, and TAO, DACHENG. “Learnable Skeleton-Aware 3D Point Cloud Sampling”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2023, 17671–17681 2.
- [XDW*23] XU, RUI, DOU, ZHIYANG, WANG, NINGNA, et al. “Globally consistent normal orientation for point clouds by regularizing the winding-number field”. *ACM Transactions on Graphics (TOG)* 42.4 (2023), 1–15 2.
- [XWD*22] XU, RUI, WANG, ZIXIONG, DOU, ZHIYANG, et al. “RFEPS: Reconstructing feature-line equipped polygonal surface”. *ACM Transactions on Graphics (TOG)* 41.6 (2022), 1–15 2.
- [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 2.
- [XZX*24] XIONG, KEZHENG, ZHENG, MAOJI, XU, QINGSHAN, et al. “Speal: Skeletal prior embedded attention learning for cross-source point cloud registration”. *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 38. 6. 2024, 6279–6287 2.

- [YDL*23] YU, ZHENGMING, DOU, ZHIYANG, LONG, XIAOXIAO, et al. “Surf-D: High-Quality Surface Generation for Arbitrary Topologies using Diffusion Models”. *arXiv preprint arXiv:2311.17050* (2023) 2.
- [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.
- [YLL*23] YANG, LEI, LIANG, YONGQING, LI, XIN, et al. “Neural parametric surfaces for shape modeling”. *arXiv preprint arXiv:2309.09911* (2023) 2.
- [YPN*23] YARIV, LIOR, PUNY, OMRI, NEVEROVA, NATALIA, et al. “Mosaic-SDF for 3D Generative Models”. *arXiv preprint arXiv:2312.09222* (2023) 2.
- [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.
- [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.
- [ZLY*23] ZHANG, CONGYI, LIN, GUYING, YANG, LEI, et al. “Surface Extraction from Neural Unsigned Distance Fields”. *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2023, 22531–22540 2.
- [ZXW*24] ZHANG, SHENG, XU, GANG, WU, HAIYAN, et al. “Medial hex-meshing: high-quality all-hexahedral mesh generation based on medial mesh”. *Engineering with Computers* (2024), 1–21 13.