

# Dynamic Skeletonization via Variational Medial Axis Sampling

QIJIA HUANG, Université de Strasbourg, France and ICube, CNRS, France

PIERRE KRAEMER, Université de Strasbourg, France and ICube, CNRS, France

SYLVAIN THERY, Université de Strasbourg, France and ICube, CNRS, France

DOMINIQUE BECHMANN, Université de Strasbourg, France and ICube, CNRS, France




Fig. 1. Gallery of skeletons obtained with our method.

We present a novel method for computing a discrete skeleton from a shape represented by a point cloud or triangle mesh. Inspired by variational shape approximation, our approach optimizes the partitioning of the input shape by minimizing an error metric defined between medial axis samples (medial spheres) and their corresponding clusters. The metric combines plane-sphere and point-sphere distance terms and the balance between these two terms enables coarse skeletons to capture the main geometric features while denser skeletons achieve a uniform distribution of medial axis samples. The sampling of the medial axis is progressively refined through an automatic process that splits medial spheres with the highest errors. Our method's efficiency also allows users to dynamically add or remove medial axis samples locally while the optimization process continuously updates the underlying partition. Skeleton connectivity is efficiently constructed by computing the dual of the optimized shape partition. Unlike previous approaches, our method does not rely on a predefined set of candidate spheres or an initial medial axis representation.

CCS Concepts: • **Computing methodologies** → **Shape analysis; Mesh models; Point-based models.**

Additional Key Words and Phrases: Medial Axis, Skeleton, Variational approach

Authors' Contact Information: Qijia Huang, Université de Strasbourg, France and ICube, CNRS, France, qijiahuang@unistra.fr; Pierre Kraemer, Université de Strasbourg, France and ICube, CNRS, France, kraemer@unistra.fr; Sylvain Thery, Université de Strasbourg, France and ICube, CNRS, France, thery@unistra.fr; Dominique Bechmann, Université de Strasbourg, France and ICube, CNRS, France, bechmann@unistra.fr.

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

SA Conference Papers '24, December 03–06, 2024, Tokyo, Japan

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1131-2/24/12

<https://doi.org/10.1145/3680528.3687678>

## ACM Reference Format:

Qijia Huang, Pierre Kraemer, Sylvain Thery, and Dominique Bechmann. 2024. Dynamic Skeletonization via Variational Medial Axis Sampling. In *SIGGRAPH Asia 2024 Conference Papers (SA Conference Papers '24)*, December 03–06, 2024, Tokyo, Japan. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3680528.3687678>

## 1 Introduction

The medial axis of a bounded shape is defined as the set of points that have at least two nearest points on the boundary of the shape, i.e. the centers of all maximally inscribed balls. Along with their associated radius, these points form an alternative representation of the original shape, the medial axis transform (MAT). Application domains of this representation include segmentation [Lin et al. 2022], deformation [Lan et al. 2020], animation [Yang et al. 2018] and hexahedral volume mesh generation [Quadros 2014; Viville et al. 2023]. In 3D, the medial axis of a surface is composed of both curve and surface structures.

Due to its complexity and sensitivity to the input shape noise, discrete approximations of the medial axis, sometimes referred to as skeletons, are commonly used [Tagliasacchi et al. 2016]. Depending on the application, some skeletons consist solely of curves, capturing essentially tubular shapes, while others can also include surface parts. The quality of a skeleton is generally evaluated through a classical trade-off between accuracy and compactness. Accuracy is typically measured by the two-sided Hausdorff distance between the original shape and the surface implicitly defined by interpolating medial spheres over the connectivity of the skeleton [Sun et al. 2015]. Compactness is simply determined by the number of medial samples.

One of the most prominent family of existing methods for shape skeletonization adopts a progressive simplification approach, starting with a fine approximation and gradually coarsening or subsampling it while minimizing error relative to the original data. Although these methods can produce accurate results, they often lack control over the compactness and distribution of the resulting samples. Additionally, the inherent operations and computation times of these methods often make it difficult for the user to intervene in the control of the result.

Inspired by successful variational approaches in geometric tasks such as shape approximation [Cohen-Steiner et al. 2004], segmentation [Lu et al. 2007] and surface reconstruction [Zhao et al. 2023], we present a novel coarse-to-fine method for the computation of a discrete skeleton of a shape represented by a point cloud or a triangle mesh. We observed that under conditions of extreme simplification, each sample of a discrete representation of a medial axis captures a distinct segment of the object’s local geometry. Motivated by this insight, our objective is to determine the optimal parameters of a given number of medial spheres so that they best represent these geometric segments. To achieve this, we develop a metric that combines a plane-sphere and a point-sphere distance terms. This metric is employed in an optimization loop to partition the input shape based on a set of medial spheres and as an objective function for updating these spheres within their cluster. Starting from a single sample, an automatic progressive refinement process adds samples according to the highest errors until a given threshold is met or a maximal spheres count is reached. The inherent efficiency allows users to interactively add or remove samples according to the user needs while the optimization process continuously updates the underlying partition. To establish the connectivity of the skeleton, we propose using the dual of the shape partition, a method fast enough to be computed on the fly. Ultimately, our method outperforms state-of-the-art methods by achieving lower reconstruction errors and offering greater control over the distribution of medial spheres, all while being fast and robust to noise.

The contributions of this paper are the following:

- a variational framework for the sampling of the medial axis of a 3D shape that does not rely on an initial representation of the medial axis or a predefined set of candidate points. Instead, our method progressively discovers the medial axis, effectively updating the positions and radii of medial spheres based on their dynamic relationship with the input shape’s clusters.
- a fast and versatile connectivity building algorithm. Combined with our highly parallelizable optimization process, our method is fast enough to enable interactive tuning of the skeleton. This capability provides enhanced flexibility and usability for a wide range of geometric modeling tasks.

## 2 Related work

*Medial Axis approximation.* The most popular methods for approximating the medial axis start with the construction of a Voronoi diagram of sample points on the surface of the shape and keeping interior points and their connectivity. The resulting structure is often too noisy and impractical for real-world applications due either

to the inherent noise in the input data or to the well-known problem of spikes generated from sliver tetrahedra [Amenta and Bern 1998] in the Delaunay triangulation of the input points. Consequently, a significant number of methods focus on simplifying the structure of the medial axis by identifying stable and meaningful parts within this medial mesh.

*Criteria-based filtering* methods eliminate the Voronoi vertices that do not meet a given criteria. Examples include the  $\lambda$ -medial axis [Chazal and Lieutier 2005] that filters samples based on their radius, or the angle-based filtering [Amenta et al. 2001; Dey and Zhao 2002; Foskey et al. 2003; Sud et al. 2005] that filters out points of the medial axis based on the angle formed by its two closest points. While simple and efficient, these approaches suffer either to preserve the topology of the original shape or its features at different scales. The *Scale Axis Transform* [Giesen et al. 2009; Miklos et al. 2010] applies a multiplicative scaling to the initial medial balls, removes those that are absorbed, and then scales back the surviving ones. A notable drawback of this method is that it can lead to changes in topology during the shrinking process.

Inspired by mesh simplification techniques like the QEM decimation method [Garland and Heckbert 1997], many approaches [Faraj et al. 2013; Li et al. 2015; Pan et al. 2019; Sun et al. 2015] have attempted to progressively simplify the medial axis approximation by performing edge collapse operation. Among these, Q-MAT [Li et al. 2015] defines a quadratic error metric that guides the simplification process, achieving highly accurate shape approximations. However, its reliance solely on local information results in a lack of control over the spheres distribution.

More recently, the Coverage Axis method [Dou et al. 2022] proposes to extract the medial axis by identifying, among a set of initial candidates, the minimum number of dilated medial spheres required to cover the entire set of the input surface points. However, since the set cover problem they address is NP-hard, it induces a very high computational cost. To keep the computation within tractable limits, the shape is sparsely sampled, thus limiting the complexity of addressable shapes. A recent update of this method [Wang et al. 2024a] mitigates this complexity, but at the expense of less control over the distribution of the selected medial spheres. Importantly, both approaches heavily rely on the quality of the initial candidate set. Inadequate candidates can lead to suboptimal results.

MATFP [Wang et al. 2022] specializes in the processing of CAD models ensuring the preservation of both internal medial axis and external mesh surface features that are pre-computed in order to constrain some samples to lie on the sharp features of the shape. A dense set of internal samples are generated from a classical Voronoi based approach and corrected as valid medial spheres using an optimization based on the geometry of their surface contact points. Careful handling of different configurations allow to maintain the features of the medial mesh and the connectivity is deduced as the dual of the restricted power diagram. Topology preservation is not guaranteed but is addressed in a following work [Wang et al. 2024b] using a volumetric representation of the shape. This method produces accurate medial axes but is not tailored to generate coarse approximations or provide dynamic control over the result.

Recent advances in deep learning has encouraged the investigation of learning-based method for skeleton extraction [Clémot




Fig. 2. **Overview of the method:** Starting from a given shape, a first cluster containing all vertices is defined and a first sphere is optimized using our update metric. After the insertion of new spheres, clusters are computed and spheres are updated in an optimization loop. The process stops when the total error meets a given threshold or a pre-defined number of samples is reached. Finally, the connectivity of the medial axis approximation is deduced from the adjacency of the clusters.

and Digne 2023; Ge et al. 2023; Lin et al. 2021; Yang et al. 2020]. In particular, Point2Skeleton [Lin et al. 2021] learns the skeletal representations by predicting and connecting skeletal points to form a mesh in an unsupervised manner. Point2MM [Ge et al. 2023] learns directly a medial mesh from point clouds. However, neither method offers interactive control over the resulting skeleton, nor do they effectively preserve the geometric features of the original shape. Neural Skeleton [Clémot and Digne 2023] trains an implicit neural representation and extracts the skeleton by combining adaptive medial point sampling with the Coverage Axis to construct connectivity, thereby sharing the same drawbacks as Coverage Axis.

*Variational approaches.* Variational approaches have been successfully applied to a variety of tasks in geometry processing, including surface approximation [Cohen-Steiner et al. 2004; Wu and Kobbelt 2005], segmentation [Yan et al. 2012], and remeshing [Yan et al. 2009]. These methods have also been extended to adapt point clouds [Skrodzki et al. 2020] for surface approximation and, more recently, for surface reconstruction [Zhao et al. 2023]. Typically, these approaches address their specific challenges by progressively optimizing a partition of the input shape. This is achieved by minimizing an error metric formulated between the clusters of the partition and various geometric proxies, such as planes, spheres, or cylinders.

To our knowledge, such an approach has not yet been used for the task of medial axis approximation. The appealing characteristics

are notably the progressive construction that allow to obtain results at different scales and the control of the result via the parameters of the optimized metric.

### 3 Method

The input of our method is a 3D shape represented either as a point cloud with normals or as a surface triangle mesh. As an output, it produces a skeleton, an approximation of its medial axis, represented as a non-manifold mesh with surface and curve parts. In the case of a point cloud, neighborhood information is required and we build a  $k$ -nearest neighbors (KNN) graph of the point cloud, using  $k = 6$  as a typical value.

Our method relies on the minimization of distances between spheres, points and planes defined at the vertices of the input shape (Sec. 3.1). The optimization of the medial spheres is performed by interleaving the partitioning (Sec. 3.2) of the input vertices – associating each medial sphere with its corresponding cluster – and the update of the medial spheres to minimize the error between each sphere and the local geometry of its cluster (Sec. 3.3). Medial spheres that exhibit substantial error are split to progressively improve the global accuracy of the skeleton (Sec. 3.4). This iterative process continues until the total error meets a given threshold or a maximum spheres count is reached. Finally, the connectivity of the skeleton is

established by computing the dual of the shape partition (Sec. 3.5). An overview of the method is illustrated in Figure 2.

*Notations.* We denote  $s = (q, r) \in \mathbb{R}^4$  the sphere of center  $q \in \mathbb{R}^3$  and radius  $r \in \mathbb{R}$ .  $\{p, n\}^\perp$  denotes the plane orthogonal to the unit vector  $n \in \mathbb{R}^3$  and passing through the point  $p \in \mathbb{R}^3$ . In the following expressions, vectors are considered as column vectors.

### 3.1 Distances definition

The distance between a sphere  $s = (q, r)$  and a point  $p$  is given by

$$d_p(s) = (\|p - q\|) - r \quad (1)$$

The distance between a sphere  $s = (q, r)$  and a plane  $\{p, n\}^\perp$  (see inset) is given by

$$d_{p,n}(s) = n^t \cdot (p - q) - r \quad (2)$$

As introduced by [Thiery et al. 2013], the square of this distance can be expressed as a quadric w.r.t. a given sphere  $s$  – namely, the spherical quadric error metric (SQEM) – that naturally compose and is defined as

$$Q_{p,n}(s) = \frac{1}{2} s^t \cdot A \cdot s - b^t \cdot s + c \quad (3)$$

where the matrix  $A$ , vector  $b$ , and scalar  $c$  are given by

$$A = 2 \begin{bmatrix} n \cdot n^t & n \\ n^t & 1 \end{bmatrix}, b = 2(n^t \cdot p) \begin{bmatrix} n \\ 1 \end{bmatrix}, c = (n^t \cdot p)^2$$

Depending on the nature of the input shape description, the functions  $D_{v_i}(s)$ , that measures the weighted squared distance between a sphere  $s$  and the vertex  $v_i$ , and  $Q_{v_i}(s)$ , that measures the weighted squared distance between a sphere  $s$  and a set of planes associated to the vertex  $v_i$ , are defined in the following way.

*Point Cloud Input.* Given a point cloud  $\mathcal{P}$ , we denote  $v_i \in \mathbb{R}^3$  a vertex of  $\mathcal{P}$ ,  $n_i \in \mathbb{R}^3$  its given or estimated normal and  $\mathcal{A}(v_i) \in \mathbb{R}$  its support area computed as proposed in [Zhao et al. 2023]. The function  $D_{v_i}$  is defined as

$$D_{v_i}(s) = \mathcal{A}(v_i) d_{v_i}(s)^2 \quad (4)$$

The diffused quadric  $Q_{v_i}$  is defined as

$$Q_{v_i}(s) = \sum_{v_j \in KNN(v_i)} \frac{\mathcal{A}(v_j)}{k} Q_{v_j, n_j} \quad (5)$$

where  $KNN(v_i)$  is the set of  $k$  nearest neighbors of  $v_i$  in  $\mathcal{P}$ .

*Mesh Input.* Given a triangle surface mesh, we denote  $v_i \in \mathbb{R}^3$  a vertex of the mesh,  $t_i$  a triangle of the mesh,  $n_i$  the normal of triangle  $t_i$  and  $\mathcal{A}(t_i)$  the area of triangle  $t_i$ . The function  $D_{v_i}$  is defined as

$$D_{v_i}(s) = \left( \sum_{t_j \in T(v_i)} \frac{\mathcal{A}(t_j)}{3} \right) d_{v_i}(s)^2 \quad (6)$$

where  $T(v_i)$  is the set of triangles incident to  $v_i$  in the mesh. The diffused quadric  $Q_{v_i}$  is defined as

$$Q_{v_i}(s) = \sum_{t_j \in T(v_i)} \frac{\mathcal{A}(t_j)}{3} Q_{v_i, n_j} \quad (7)$$

### 3.2 Shape partitioning

Our method starts with a single medial sphere, derived by considering the entire set of vertices as one cluster and optimizing our metric (Sec 3.3). As new medial spheres are subsequently added, it becomes necessary to compute a new partition of the input vertices. To achieve this, each vertex  $v_i$  of the shape is iteratively processed and assigned to the cluster of the medial sphere  $m_j = (q_j, r_j)$  that minimizes a cost. The cost associated with adding vertex  $v_i$  to the cluster of sphere  $m_j$  is defined as:

$$E_{v_i}(m_j) = Q_{v_i}(m_j) + \lambda D_{v_i}(m_j) \quad (8)$$

This cost function combines a SQEM term  $Q_{v_i}$  and a Euclidean distance term  $D_{v_i}$  weighted by a coefficient  $\lambda$ . Using only the SQEM term, which measures plane-sphere distances, would result in large anisotropic clusters with several components and noisy borders, as this distance remains low in flat or tubular regions, even for vertices far from the medial sphere. Adding a point-sphere distance term allows to define more compact regions with clean borders. This is illustrated in Figure 3 with a shape composed essentially of flat and cylinder parts. During the optimization process, with the progressive addition of medial spheres, the SQEM error tends to decrease rapidly, giving more importance to the Euclidean term, which leads to more compact and isotropic clusters.




Fig. 3. **The effect of  $\lambda$  in partitioning:** For a same set of spheres (left), small values of  $\lambda$  results in clusters with noisy boundaries, leading to poor adjacency quality for the connectivity, while higher values of  $\lambda$  lead to more compact and regular clusters which in turn yield more meaningful connectivity.

### 3.3 Spheres updating

The next step involves adjusting each medial sphere  $m_i \in \mathcal{M}$  to match the local geometry of its cluster  $C_i$ . For each cluster  $C_i$ , the optimal sphere that most accurately reflects the local geometry is identified by minimizing the total squared distances to all planes defined at the vertices  $v_j$  within the cluster. Using the quadrics associated to each vertex  $v_j$ , the distance metric is defined as:

$$E_{SQEM}(C_i) = \sum_{v_j \in C_i} Q_{v_j}(m_i) \quad (9)$$

However, minimizing this function is ill-posed in open cylindrical or flat regions as an infinite number of spheres can achieve minimal distances in such configurations. For example, any of the blue spheres depicted in Figure 4 would equally fit the constraints carried by the planes defined on the surface of the cluster.




Fig. 4. Within an open cylindrical cluster, an infinite number of spheres (blue) would minimize the SQEM energy. Adding a small-weight Euclidean term resolves this ambiguity without overstretching the spheres (yellow).

To address this issue, we reintroduce the Euclidean distance term into our optimization equation. The following term computes the sum of the Euclidean distances from the medial sphere  $m_i = (q_i, r_i)$  to each of the vertices of the cluster  $C_i$ :

$$E_{\text{euclidean}}(C_i) = \sum_{v_j \in C_i} D_{v_j}(m_i) \quad (10)$$

The optimal sphere  $m_i^* = (q_i^*, r_i^*)$  for the cluster  $C_i$  is then derived by minimizing the combined objective function:

$$(q_i^*, r_i^*) = \arg \min_{q_i, r_i} (E_{\text{SQEM}}(C_i) + \lambda E_{\text{euclidean}}(C_i)) \quad (11)$$

where the Euclidean distance term is weighted by  $\lambda$ . This optimization problem is inherently non-linear, and the optimal sphere can be efficiently computed using Gauss-Newton method, with the current sphere configuration serving as the initial guess.

At the beginning of the optimization process, the number of medial spheres is small and most clusters are well-defined w.r.t. the SQEM energy minimization. At this stage, this term is crucial and naturally drives the spheres towards the main features of the shape like extremities of cylindrical regions or corners of flat areas. Once these critical features are adequately captured, new spheres are more likely to be added in the middle of flat or tubular regions. The Euclidean term then plays a more significant role, promoting an even distribution of spheres within the intermediate spaces. Figure 5 illustrates the impact of the parameter  $\lambda$  during the optimization process on a simple example. When  $\lambda = 0$  (only the SQEM term is used), the two spheres perfectly fit the extremities of the shape, but any additional sphere conflicts with either one of them. The introduction of the Euclidean term slightly pushes the spheres off the extremities and beyond the surface of the shape but allows additional spheres to be distributed evenly. Large values of  $\lambda$  tends to favor the uniform distribution of spheres at the expense of the fidelity to the features of the shape. As will be exposed in the results, we consistently use a value of  $\lambda = 0.2$  which yields a good compromise.

*Error computation.* After an update step, the error of each cluster  $C_i$  is computed by summing the errors of each vertex  $v_i \in C_i$  w.r.t. the medial sphere of the cluster  $m_i$  following Equation 8, normalized

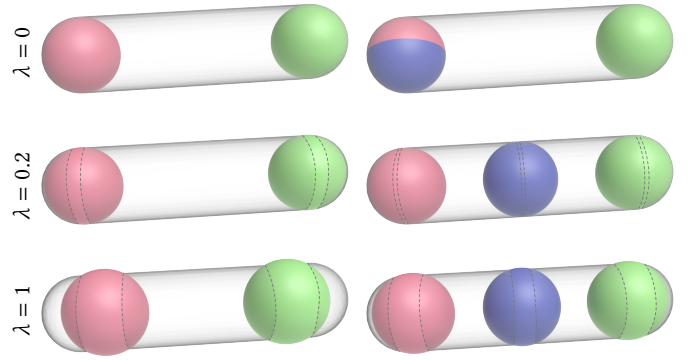


Fig. 5. **The effect of  $\lambda$  in updating:** Pure SQEM ( $\lambda = 0$ ) effectively fits spheres within the features of the shape. However, additional spheres conflict with the existing ones. Introducing a Euclidean term addresses this instability, evenly distributing the space between samples, but as  $\lambda$  increases, spheres tend to stretch and extend beyond the surface boundary.

by the total area of the cluster  $\mathcal{A}(C_i)$ :


$$E(C_i) = \frac{1}{\mathcal{A}(C_i)} \sum_{v_j \in C_i} E_{v_j}(m_i) \quad (12)$$

Each optimization step (partitioning and updating) decreases the total error over all clusters. The process terminates when the reduction in error is less than a predefined small threshold  $\epsilon$ .

*Sphere projection.* As already illustrated in Fig 5, the minimization of this energy function does not provide any guarantee that the obtained spheres remain within the boundaries of the shape. To address a similar issue, where the optimization results are not well restricted within the desired domain, remeshing [Yan et al. 2009] or reconstruction [Zhao et al. 2023] variational approaches project the updated seeds or generators within the object space between the iterations.

We devise such a *correction* step that projects back a sphere onto the medial axis (see inset). Given a sphere  $s = (q^*, r^*)$ , we first identify the point  $p$  on the shape that is closest to the sphere center  $q^*$ . A bounding volume hierarchy for surface inputs or a KD-tree for point cloud inputs is used to speed-up the process. Next, we calculate the direction vector  $d = \frac{q^* - p}{\|q^* - p\|}$  which aligns with the gradient of the distance function. Finally, we apply the shrinking ball algorithm [Ma et al. 2012] from point  $p$  in the direction  $d$  which computes the corrected medial sphere  $m = (\hat{q}, \hat{r})$ .

While the correction step temporarily interrupts the convergence of the optimization process locally, it does not compromise the overall convergence of the algorithm, as demonstrated in Fig. 11. Applying the correction step after each optimization iteration results in a similar final error as deferring the correction until the algorithm's completion.



### 3.4 Spheres splitting

In order to improve the quality of the medial axis approximation, new samples can be added via spheres splitting. Clusters are sorted by their error and considered in turn in descending order. A cluster is considered for splitting if its error is above a given threshold and if its adjacent clusters have not been considered for splitting (efficient construction of the adjacency between clusters is addressed in the next section 3.5). For each cluster selected for splitting, we identify the vertex that carries the highest error. These vertices are used as seeds to add new medial spheres into  $\mathcal{M}$ . Initial configuration of these new spheres is obtained using the shrinking ball algorithm [Ma et al. 2012].

When the optimization has converged and no cluster has an error greater than a given threshold or a prescribed maximal number of spheres is met, the global optimization process is stopped. Along with  $\lambda$ , this error threshold or maximal number of spheres is the second parameter that allows to control the result, a lower threshold leading to finer and more accurate skeletons.

### 3.5 Connectivity building

After convergence of the optimization of the set of medial spheres  $\mathcal{M}$ , the skeletal structure is built by establishing a connectivity between the sphere centers. Several strategies have been proposed in the literature. Dou et al. [Dou et al. 2022] suggested to embed the selected spheres into a medial axis represented by a Voronoi diagram and simplify it using the Q-MAT algorithm [Li et al. 2015]. While being effective in maintaining topological consistency and minimal reconstruction error, this method is computationally intensive (c.f. Table 2). Other methods such as MATFP [Wang et al. 2022], which is inspired by the power shape [Amenta et al. 2001], builds the connectivity by computing the restricted regular triangulation, which is the dual of restricted power diagram of the set of spheres.

Aiming at efficiency and simplicity, we propose a method similar to that of MATFP [Wang et al. 2022] that enables fast connectivity construction. We observe that as the number of clusters increases, the SQEM error decreases and the Euclidean distance term becomes predominant. Consequently, the emerging clusters are closely aligned with the power diagram of the medial spheres restricted to the set of vertices of the input shape. This motivates us to construct the connectivity of the medial spheres by computing the dual of their associated clusters connectivity. More specifically, an edge is inserted between two sphere centers if their respective clusters are adjacent, and a face is formed between three sphere centers if their clusters are mutually adjacent. Adjacency between clusters can be easily obtained in meshes or point clouds by going through the edges of the mesh or the KNN-graph.

## 4 Results

Our algorithm has been implemented in C++ with the CGoGN library for meshes [Kraemer et al. 2020] and the Eigen library for linear algebra operations [Guennebaud et al. 2010]. All tests were performed on a computer equipped with an Intel(R) Core(TM) i7-12700k CPU at 3.60 GHz and 32 GB of RAM. We normalized the size of all models to a [0,1] range.

We conducted comparisons primarily against Coverage Axis (CA) [Dou et al. 2022], Coverage Axis++ (CA+) [Wang et al. 2024a], and Q-MAT [Li et al. 2015], as these methods are also able to generate skeletons with a specific number of spheres. For fairness, we used CA to determine the number of spheres for each model, and results of both three methods are post-processed with a correction step using the shrinking ball algorithm. Following [Dou et al. 2022; Wang et al. 2024a], we set the dilation constant  $\delta_r$  to 0.2, sampled 1500 vertices uniformly from the model surface, and generated 20000 random vertices within the model. The Set Cover problem is then solved using the HIGHS library [Huangfu and Hall 2018], and the number of selected vertices determines the target number of spheres. To establish the connectivity for CA and CA+, we embedded the selected spheres in Q-MAT and simplified it until only the selected spheres remained. This process was achieved using a program provided by the authors. For our method, we consistently set  $\lambda = 0.2$  throughout the experiments.

*Reconstruction Fidelity.* One of the most important properties of the medial axis is its ability to accurately reconstruct the original object. The reconstructed shape is obtained by interpolating the medial spheres along the edges and faces of the medial mesh [Li et al. 2015; Sun et al. 2015]. We evaluated the reconstruction quality using the two-sided Hausdorff Distance (HD), denoted as  $\epsilon$ . Results are shown in Table 1. We can observe that, although Q-MAT is designed for highly accurate shape approximations, our results are superior in most cases. This discrepancy arises because Q-MAT, as a progressive simplification method, accumulates errors during extreme simplifications, since it considers only local information at each step. Conversely, our approach integrates global information during the partitioning process, thereby minimizing error accumulation and resulting in more precise outcomes.

*Time Efficiency.* Table 2 presents quantitative data on the efficiency of our method. We tracked execution time across various stages, including initial medial axis computation (CA, CA+, Q-MAT),

Table 1. Reconstruction error comparison with Coverage Axis, Coverage Axis++ and Q-MAT.  $|M|$ : The number of skeleton points,  $\epsilon$ : Two-sided Hausdorff distance

Method		Coverage Axis	Coverage Axis++	Q-MAT	Ours
Model( $ V $ )	$ M $	$\epsilon$	$\epsilon$	$\epsilon$	$\epsilon$
Armadillo (36725)	83	3.438%	3.905%	3.860%	<b>2.735%</b>
Fertility (17827)	98	2.364%	3.386%	3.984%	<b>1.630%</b>
Chair (10500)	117	1.892%	2.412%	1.606%	<b>1.123%</b>
Vase (14859)	117	2.467%	2.599%	2.847%	<b>1.890%</b>
Elephant(24955)	117	2.074%	2.733%	2.514%	<b>1.741%</b>
Bug(8640)	75	2.338%	2.429%	2.117%	<b>1.001%</b>
Seahorse(20494)	76	3.195%	3.393%	5.088%	<b>2.967%</b>
Dove(5519)	76	2.754%	2.885%	2.186%	<b>0.952%</b>
Vessel (49698)	99	3.054%	3.033%	<b>0.818%</b>	1.716%
Pinion (10369)	230	3.393%	3.319%	2.194%	<b>2.127%</b>
Elk (24013)	124	2.478%	2.258%	1.697%	<b>0.944%</b>
Neptune (14814)	106	4.063%	3.199%	3.329%	<b>2.374%</b>
Dolphin (15100)	51	2.043%	2.446%	1.983%	<b>1.296%</b>
Santa (10241)	90	2.011%	2.372%	2.158%	<b>1.475%</b>
Bear (10141)	45	3.285%	3.391%	2.204%	<b>1.659%</b>
Plane (7651)	86	2.151%	2.945%	1.643%	<b>1.106%</b>
Spider (11051)	107	1.692%	1.604%	0.901%	<b>0.495%</b>
Venus (10760)	43	<b>2.162%</b>	2.701%	3.768%	2.929%

Table 2. Time efficiency comparison with Coverage Axis, Coverage Axis++ and Q-MAT. Times are measured in seconds.

Method	Model( V )	M	MA init	Coverage Axis			Coverage Axis++			Q-MAT		Ours
				Selection	Connectivity	Total time	Selection	Connectivity	Total time	Simplification	Total time	
Armadillo (36725)	83	4.528	3.6	10.476	18.604	0.93	10.753	16.211	10.323	14.851	<b>2.179</b>	
Fertility (17827)	98	1.811	24.29	5.012	31.113	1.106	4.855	7.772	4.632	6.443	<b>1.46</b>	
Chair (10500)	117	1.091	509.4	2.481	512.972	1.31	2.347	4.748	2.091	3.182	<b>0.971</b>	
Vase (14859)	117	1.512	1.3	4.095	6.907	1.3	4.157	6.969	4.256	5.768	<b>0.9</b>	
Elephant (24955)	117	2.703	3.84	6.933	13.476	1.3	7.359	11.362	6.554	9.257	<b>2.226</b>	
Bug (8640)	75	0.91	5.1	1.96	7.97	0.9	1.991	3.801	1.848	2.758	<b>0.429</b>	
Seahorse (20494)	76	2.019	18.85	5.494	26.363	0.9	5.511	8.43	5.074	7.093	<b>1.035</b>	
Dove (5519)	76	0.53	>1000	1.017	>1000	0.88	1.186	2.596	1.029	1.559	<b>0.42</b>	
Vessel (49698)	99	5.337	1.2	11.849	18.386	1.127	12.587	19.051	11.185	16.522	<b>2.495</b>	
Pinion (10369)	230	1.02	2.3	2.295	5.615	2.79	2.289	6.099	2.439	3.459	<b>2.584</b>	
Elk (24013)	124	2.538	2.62	6.108	11.266	3.485	6.228	12.251	5.746	8.284	<b>1.94</b>	
Neptune (14814)	106	1.148	5.6	3.478	10.226	1.196	3.562	5.9063	3.221	4.369	<b>1.017</b>	
Dolphin (15100)	51	1.761	4.47	4.568	10.799	0.7	4.564	7.025	4.31	6.071	<b>0.751</b>	
Santa (10241)	90	1.07	3.89	2.628	7.588	1.04	2.768	4.878	2.529	3.599	<b>0.598</b>	
Bear (10141)	45	1.1	2.31	2.76	6.17	0.61	2.803	4.513	2.68	3.78	<b>0.34</b>	
Plane (7651)	86	0.751	262.1	1.999	264.85	0.973	2.039	3.763	1.51	2.262	<b>0.529</b>	
Spider (11051)	107	1.176	1.34	2.733	5.249	1.158	2.664	4.998	2.536	3.712	<b>0.843</b>	
Venus (10760)	43	1.11	3.02	3.023	7.154	0.597	2.973	4.681	2.659	3.77	<b>0.343</b>	

selection of the medial samples (CA, CA+), simplification (Q-MAT), and connectivity construction (CA, CA+). Since our method constructs connectivity at each iteration (to prevent adjacent spheres splitting), we do not provide detailed timing for each stage of the algorithm. Owing to the local nature of our sphere updating process, which is highly amenable to parallelization, our method generally outperforms other approaches by an order of magnitude in terms of efficiency.

*Visual comparison.* Figure 9 shows the results obtained with these different methods for different shapes. For each model, we show the obtained medial axis approximation, the reconstructed shape and a color-coded representation of the reconstruction error. Figure 10 shows the medial axis approximation obtained for a shape at several resolution levels. As observed, both CA and CA+ struggles to generate a regular distribution of spheres and fails to preserve geometric features with randomly generated candidates. Similarly, Q-MAT, focused on local metric changes, does not maintain overall distribution or features effectively. In contrast, our method achieves a more uniform sphere distribution while better preserving the shape’s key geometric features.

*Robustness to noise.* We assessed the robustness of our method by adding noise to the input shape vertices position with different magnitudes. Figure 6 shows how our method is still able to produce good results up to a high level of noise (expressed in % of the bounding box diagonal). Note that starting from 3% of noise, the correction step was not applied anymore, as the surface itself becomes fuzzy and the projection process prevents the optimization from being stable.

*Study of the metric.* Figure 8 shows the role of the different terms of the metric on a simple example shape. Using only the Euclidean term, optimized spheres are not constrained in the shape. The projection step helps to put them back in the shape but the geometric features are not well captured. Using only the SQEM term, the underlying clusters get too extended which leads to spheres that reach

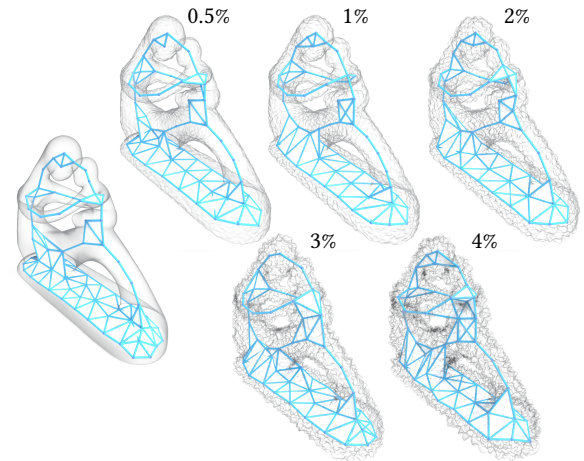


Fig. 6. **Robustness to noise:** Our method still produces good results even with a fairly high level of noise in the input data.

beyond the shape and irrelevant connectivity. The proposed combined metric leads to more compact clusters, better fitted spheres and proper connectivity. The projection step merely corrects the optimized spheres to ensure they are actual medial axis samples. Figure 7 shows the effect of the  $\lambda$  coefficient, which controls the influence of the Euclidean term, for a more complex shape than shown in Sec. 3. A small value leads to a perfect fit of the medial samples in the features of the shape, but as the number of samples increases, they tend to conflict in these areas of the shape. A high value promotes regularly distributed samples which leads to a poor capture of the shape features at coarse scales. An intermediate value yields a good compromise between those behaviors.

*Interactive modification of skeleton.* An appealing aspect of our method is the ability to interactively edit the skeleton. The user can remove or split the hovered medial sample – even without having the spheres to be actually displayed, while the optimization loop




Fig. 7. **Effect of the  $\lambda$  parameter.** Medial spheres and skeleton obtained for a plane model with different values of  $\lambda$  and different number of medial samples.

continues to update the medial mesh in the background. Figure 12 shows an interactively adapted skeleton that features both high and low samples densities. This possibility is also illustrated in the accompanying video.

## 5 Limitations and future work

*Stability.* Similar to other variational partition algorithms, our method lacks a theoretical proof of global convergence and optimality. In some configurations, medial spheres may oscillate between two positions due to the discrete clustering based on vertices. This issue is more likely to happen if the user aims at a dense medial axis approximation in which some clusters may become too small to yield a stable solution. Upsampling of the object can solve this problem in most cases. However, in the surface setting, an alternative solution could involve clustering based on continuous surfaces rather than solely depending on the vertices. While this approach might increase computational costs, efficient algorithms capable of handling such tasks have already been developed [Yan et al. 2009].

*Topology preservation and connectivity.* The medial axis should consist of curves and surfaces and be homotopic to the input shape. At a very coarse resolution, the topological genus of our results can be different to that of the shape. This can be of interest to yield segmentations at different scales. However, we should also provide a refinement rule to guarantee the alignment of the topology of the medial mesh with that of the input shape. In addition, our strategy for connectivity construction that relies solely on cluster adjacency is fast but can also lead to overlapping triangles or extra cavities in the medial mesh. Ultimately, as already discussed in [Wang et al. 2022], even with a fine medial axis approximation and surface clusters that exactly correspond to the restricted power diagram of the medial spheres, this problem cannot be solved properly without the

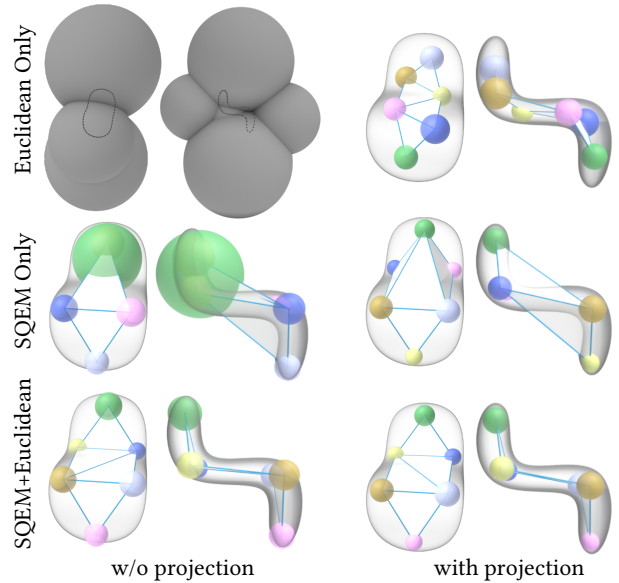


Fig. 8. **Ablation study of the metric.** Medial spheres and skeleton obtained for 6 samples on a synthetic object in front and side view under different configurations of optimization: with or w/o Euclidean term, with or w/o SQEM term, with or w/o sphere projection.

computation of a volumetric restricted power diagram. This has been done in [Wang et al. 2024b] that checks the equivalence of the connected components and Euler characteristics of the medial cells and their dual volumetric restricted elements. Pathological configurations can be detected and guide the insertion of new samples. However, the cost induced by such a scenario would hinder



our goal of providing a controllable and fast solution that supports interactive manipulation of the output. If necessary, though, such a procedure could be considered as post-processing.

*Further improvements and applications.* Several features could improve the control over the result in an interactive modeling session. Medial sample freezing could lock a sample at a given location and prevent the optimization process to further move it. A space varying density function could be designed to guide the refinement to different resolutions in different areas of the object. We also plan to study the use of our method with different kinds of input data such as binary discrete images or incomplete data.

## Acknowledgments

This research was funded by the Grand Est Region through grant 22-DOC-067 (Def-RV project) and the French government's National Research Agency as part of the Investments for the Future Program (PIA), under grant ANR-22-CE38-0010-01 (POSTURES project).

## References

- Nina Amenta and Marshall Bern. 1998. Surface reconstruction by Voronoi filtering. In *Proceedings of the Fourteenth Annual Symposium on Computational Geometry* (Minneapolis, Minnesota, USA) (SCG '98). Association for Computing Machinery, New York, NY, USA, 39–48. <https://doi.org/10.1145/276884.276889>
- Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. 2001. The power crust. In *Proceedings of the sixth ACM symposium on Solid modeling and applications*. 249–266.
- Frédéric Chazal and André Lieutier. 2005. The “ $\lambda$ -medial axis”. *Graphical models* 67, 4 (2005), 304–331.
- Mattéo Clémot and Julie Digne. 2023. Neural skeleton: Implicit neural representation away from the surface. *Computers & Graphics* 114 (2023), 368–378.
- David Cohen-Steiner, Pierre Alliez, and Mathieu Desbrun. 2004. Variational shape approximation. In *ACM SIGGRAPH 2004 Papers*. 905–914.
- Tamal K Dey and Wulue Zhao. 2002. Approximate medial axis as a voronoi subcomplex. In *Proceedings of the seventh ACM symposium on Solid modeling and applications*. 356–366.
- Zhiyang Dou, Cheng Lin, Rui Xu, Lei Yang, Shiqing Xin, Taku Komura, and Wenping Wang. 2022. Coverage Axis: Inner Point Selection for 3D Shape Skeletonization. *Computer Graphics Forum* 41, 2 (2022), 419–432. <https://doi.org/10.1111/cgf.14484>
- Noura Faraj, Jean-Marc Thiery, and Tamy Boubekeur. 2013. Progressive medial axis filtration. In *SIGGRAPH Asia 2013 Technical Briefs*. 1–4.
- Mark Foskey, Ming C Lin, and Dinesh Manocha. 2003. Efficient computation of a simplified medial axis. In *Proceedings of the eighth ACM symposium on Solid modeling and applications*. 96–107.
- Michael Garland and Paul S. Heckbert. 1997. Surface simplification using quadric error metrics. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97)*. ACM Press/Addison-Wesley Publishing Co., USA, 209–216. <https://doi.org/10.1145/258734.258849>
- Mengyuan Ge, Junfeng Yao, Baorong Yang, Ningna Wang, Zhonggui Chen, and Xiaohu Guo. 2023. Point2MM: Learning medial mesh from point clouds. *Computers & Graphics* 115 (2023), 511–521.
- Joachim Giesen, Balint Miklos, Mark Pauly, and Camille Wormser. 2009. The scale axis transform. In *Proceedings of the twenty-fifth annual symposium on Computational geometry*. 106–115.
- Gaël Guennebaud, Benoît Jacob, et al. 2010. Eigen v3. <http://eigen.tuxfamily.org>.
- Q. Huangfu and J.A.J. Hall. 2018. Parallelizing the dual revised simplex method. *Mathematical Programming Computation* 10, 1 (2018), 119–142. <https://doi.org/10.1007/s12532-017-0130-5>
- Pierre Kraemer, Sylvain Thery, et al. 2020. CGoGN. [https://github.com/cgogn/CGoGN\\_3](https://github.com/cgogn/CGoGN_3).
- Lei Lan, Ran Luo, Marco Fratarcangeli, Weiwei Xu, Huamin Wang, Xiaohu Guo, Junfeng Yao, and Yin Yang. 2020. Medial Elastics: Efficient and Collision-Ready Deformation via Medial Axis Transform. *ACM Trans. Graph.* 39, 3, Article 20 (apr 2020), 17 pages. <https://doi.org/10.1145/3384515>
- Pan Li, Bin Wang, Feng Sun, Xiaohu Guo, Caiming Zhang, and Wenping Wang. 2015. Q-mat: Computing medial axis transform by quadratic error minimization. *ACM Transactions on Graphics (TOG)* 35, 1 (2015), 1–16.
- Cheng Lin, Changjian Li, Yuan Liu, Nenglu Chen, Yi-King Choi, and Wenping Wang. 2021. Point2skeleton: Learning skeletal representations from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 4277–4286.
- Cheng Lin, Lingjie Liu, Changjian Li, Leif Kobbelt, Bin Wang, Shiqing Xin, and Wenping Wang. 2022. SEG-MAT: 3D Shape Segmentation Using Medial Axis Transform. *IEEE transactions on visualization and computer graphics* 28, 6 (2022), 2430–2444.
- Lin Lu, Yi-King Choi, Wenping Wang, and Myung-Soo Kim. 2007. Variational 3D Shape Segmentation for Bounding Volume Computation. *Computer Graphics Forum* (2007). <https://doi.org/10.1111/j.1467-8659.2007.01055.x>
- Jae-hwan Ma, Sang Won Bae, and Sunghee Choi. 2012. 3D medial axis point approximation using nearest neighbors and the normal field. *Vis. Comput.* 28, 1 (jan 2012), 7–19. <https://doi.org/10.1007/s00371-011-0594-7>
- Balint Miklos, Joachim Giesen, and Mark Pauly. 2010. Discrete scale axis representations for 3D geometry. In *ACM SIGGRAPH 2010 papers*. 1–10.
- Yiling Pan, Bin Wang, Xiaohu Guo, Hua Zeng, Yuexin Ma, and Wenping Wang. 2019. Q-MAT+: An error-controllable and feature-sensitive simplification algorithm for medial axis transform. *Computer Aided Geometric Design* 71 (2019), 16–29.
- William Roshan Quadros. 2014. LayTracks3D: A New Approach to Meshing General Solids using Medial Axis Transform. *Procedia Engineering* 82 (2014), 72–87. <https://doi.org/10.1016/j.proeng.2014.10.374> 23rd International Meshing Roundtable (IMR23).
- Martin Skrodzki, Eric Zimmermann, and Konrad Polthier. 2020. Variational shape approximation of point set surfaces. *Computer Aided Geometric Design* 80 (2020), 101875. <https://doi.org/10.1016/j.cagd.2020.101875>
- Avneesh Sud, Mark Foskey, and Dinesh Manocha. 2005. Homotopy-preserving medial axis simplification. In *Proceedings of the 2005 ACM Symposium on Solid and Physical Modeling* (Cambridge, Massachusetts) (SPM '05). Association for Computing Machinery, New York, NY, USA, 39–50. <https://doi.org/10.1145/1060244.1060250>
- Feng Sun, Yi-King Choi, Yizhou Yu, and Wenping Wang. 2015. Medial meshes—a compact and accurate representation of medial axis transform. *IEEE transactions on visualization and computer graphics* 22, 3 (2015), 1278–1290.
- Andrea Tagliasacchi, Thomas Delame, Michela Spagnuolo, Nina Amenta, and Alexandru Telea. 2016. 3D Skeletons: A State-of-the-Art Report. *Computer Graphics Forum* 35, 2 (2016), 573–597. <https://doi.org/10.1111/cgf.12865>
- Jean-Marc Thiery, Émilie Guy, and Tamy Boubekeur. 2013. Sphere-meshes: Shape approximation using spherical quadric error metrics. *ACM Transactions on Graphics (TOG)* 32, 6 (2013), 1–12.
- P. Viville, P. Kraemer, and D. Bechmann. 2023. Meso-Skeleton Guided Hexahedral Mesh Design. *Computer Graphics Forum* 42, 7 (2023), e14932. <https://doi.org/10.1111/cgf.14932> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14932>
- Ningna Wang, Hui Huang, Shibo Song, Bin Wang, Wenping Wang, and Xiaohu Guo. 2024b. MAT-Topo: Topology-preserving Medial Axis Transform with Restricted Power Diagram. *arXiv preprint arXiv:2403.18761* (2024).
- Ningna Wang, Bin Wang, Wenping Wang, and Xiaohu Guo. 2022. Computing medial axis transform with feature preservation via restricted power diagram. *ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–18.
- Zimeng Wang, Zhiyang Dou, Rui Xu, Cheng Lin, Yuan Liu, Xiaoxiao Long, Shiqing Xin, Taku Komura, Xiaoming Yuan, and Wenping Wang. 2024a. Coverage Axis++: Efficient Inner Point Selection for 3D Shape Skeletonization. arXiv:2401.12946 [cs.CV]
- Jianhua Wu and Leif Kobbelt. 2005. Structure Recovery via Hybrid Variational Surface Approximation. In *Comput. Graph. Forum*, Vol. 24. 277–284.
- Dong-Ming Yan, Wenping Wang, Yang Liu, and Zhouwang Yang. 2012. Variational mesh segmentation via quadric surface fitting. *Comput. Aided Des.* 44 (2012), 1072–1082. <https://api.semanticscholar.org/CorpusID:16945984>
- Dong-Ming Yan, Bruno Lévy, Yang Liu, Feng Sun, and Wenping Wang. 2009. Isotropic remeshing with fast and exact computation of Restricted Voronoi Diagram. In *Proceedings of the Symposium on Geometry Processing* (Berlin, Germany) (SGP '09). Eurographics Association, Goslar, DEU, 1445–1454.
- Baorong Yang, Junfeng Yao, and Xiaohu Guo. 2018. DMAT: Deformable Medial Axis Transform for Animated Mesh Approximation. *Computer Graphics Forum* 37, 7 (2018), 301–311. <https://doi.org/10.1111/cgf.13569> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13569>
- Baorong Yang, Junfeng Yao, Bin Wang, Jianwei Hu, Yiling Pan, Tianxiang Pan, Wenping Wang, and Xiaohu Guo. 2020. P2MAT-NET: Learning medial axis transform from sparse point clouds. *Computer Aided Geometric Design* 80 (2020), 101874.
- Tong Zhao, Laurent Busé, David Cohen-Steiner, Tamy Boubekeur, Jean-Marc Thiery, and Pierre Alliez. 2023. Variational Shape Reconstruction via Quadric Error Metrics. In *ACM SIGGRAPH 2023 Conference Proceedings* (Los Angeles, CA, USA.) (SIGGRAPH '23). Association for Computing Machinery, New York, NY, USA, Article 45, 10 pages. <https://doi.org/10.1145/3588432.3591529>

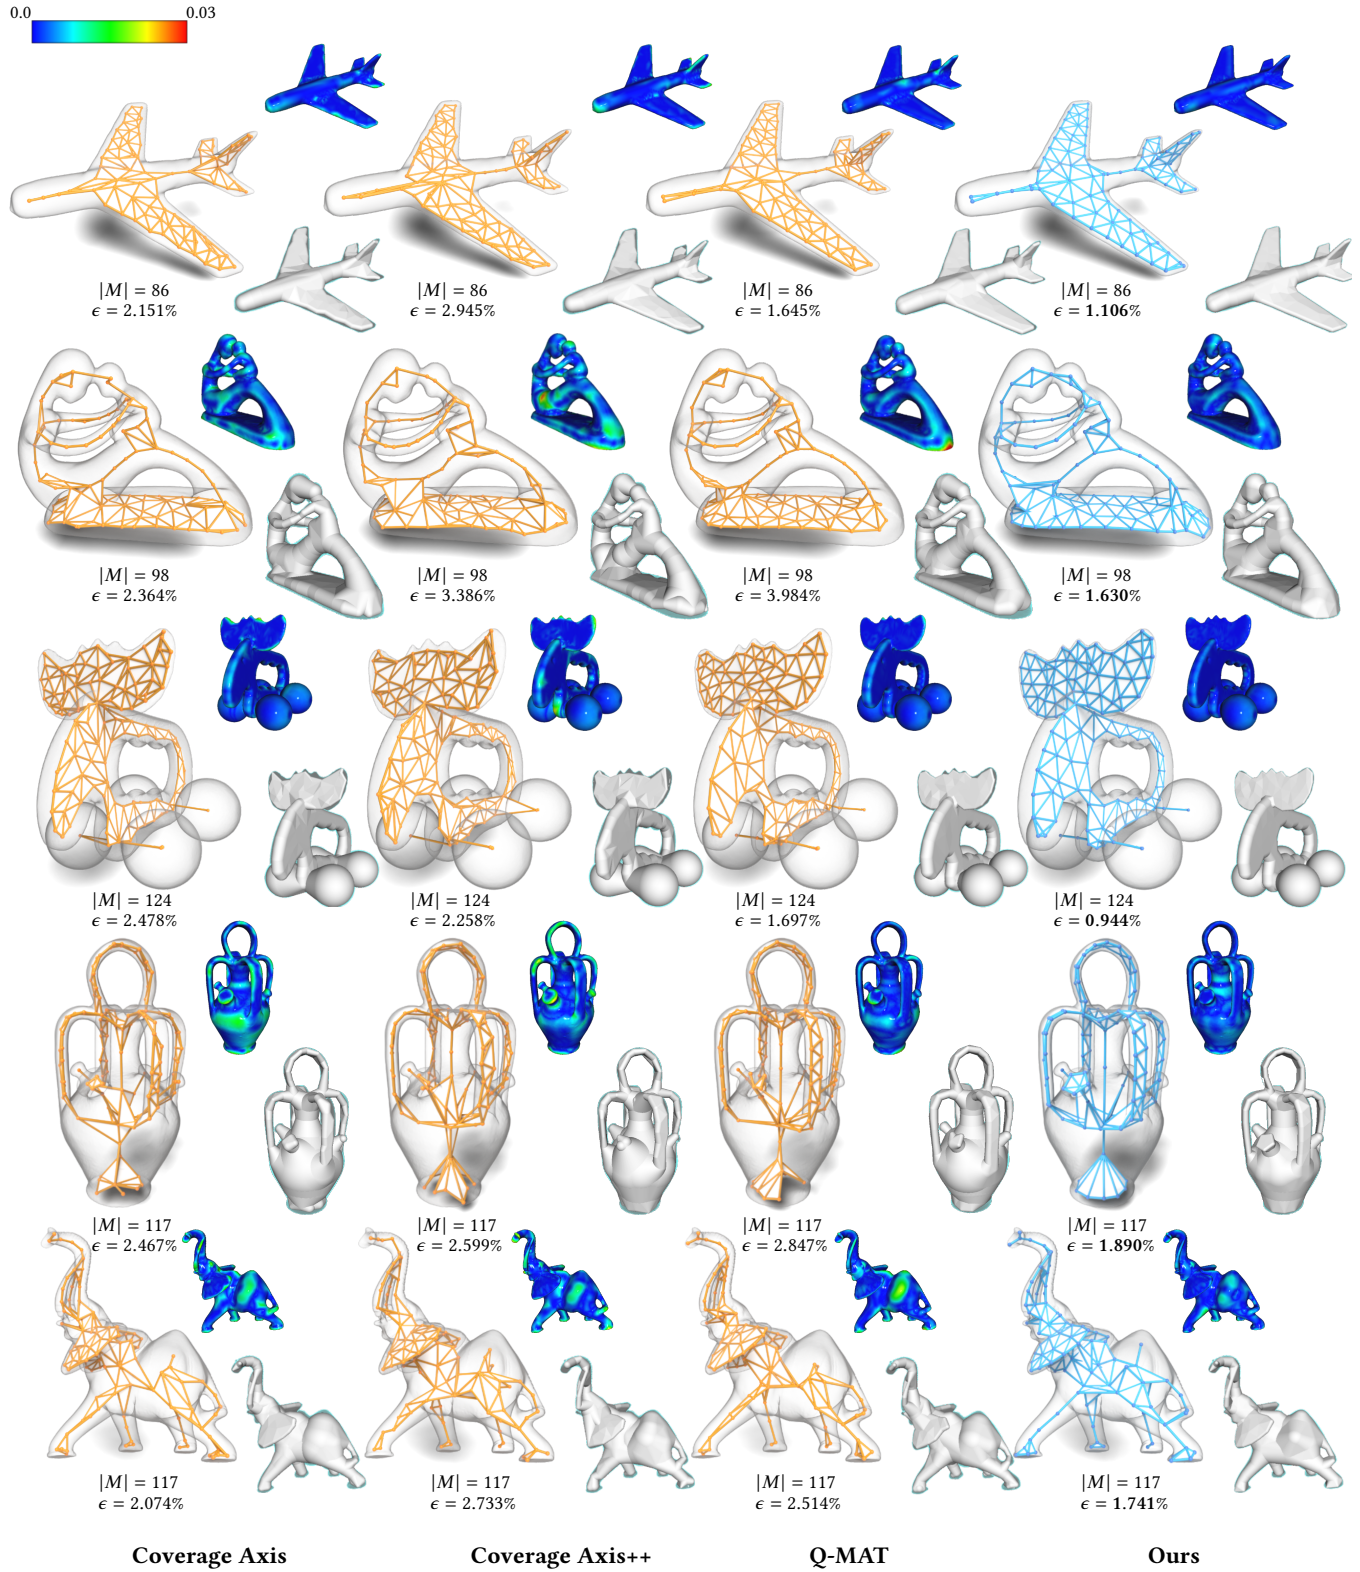


Fig. 9. Comparison of the results obtained with Coverage Axis, Coverage Axis++, Q-MAT and our method on different shapes (plane, fertility, elk, vase and elephant). Number of medial spheres  $|M|$  and reconstruction error  $\epsilon$  are displayed for each model and method.

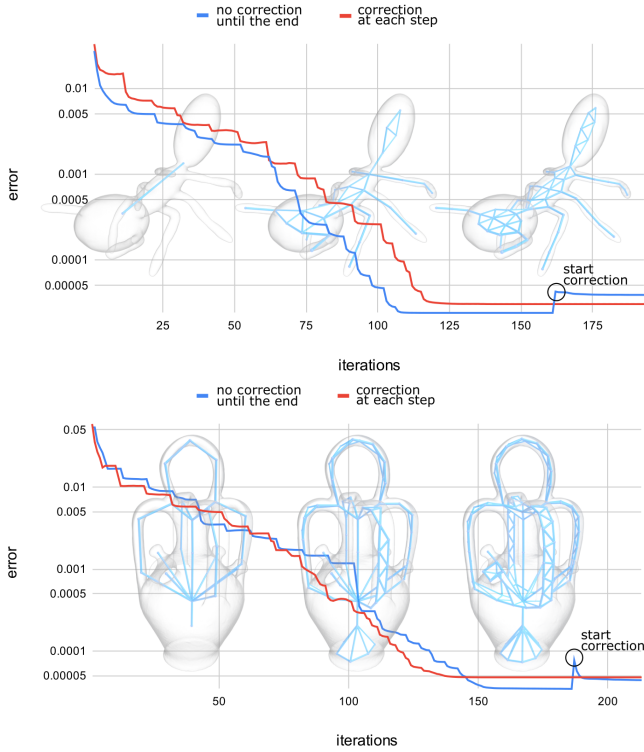


Fig. 11. **Error convergence** (on a logarithmic scale) is compared for two strategies: sphere corrections applied continuously during the optimization process (red), or only at the end (blue). The steep steps in the decrease are due to the insertion of new medial spheres. When no correction is applied during the optimization process, we activate the correction step at the end to ensure the resulting spheres align with the medial axis, which slightly increases the error. Nevertheless, the final error levels are comparable to those observed when corrections are applied continuously.

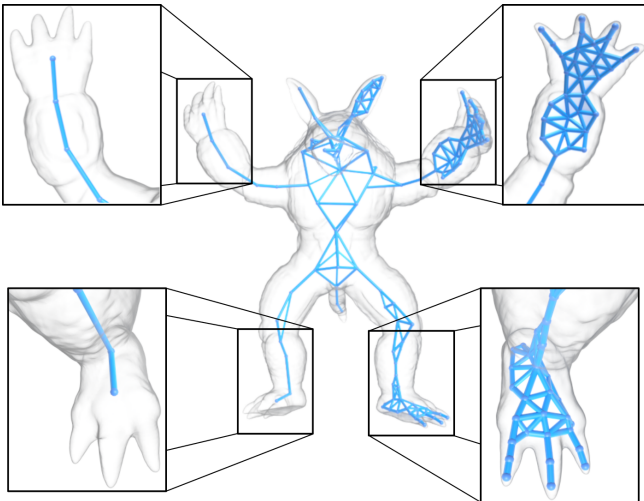


Fig. 12. **Interactive edition** via medial spheres addition and deletion allows to generate skeletons with fine and coarse parts according to the user needs.

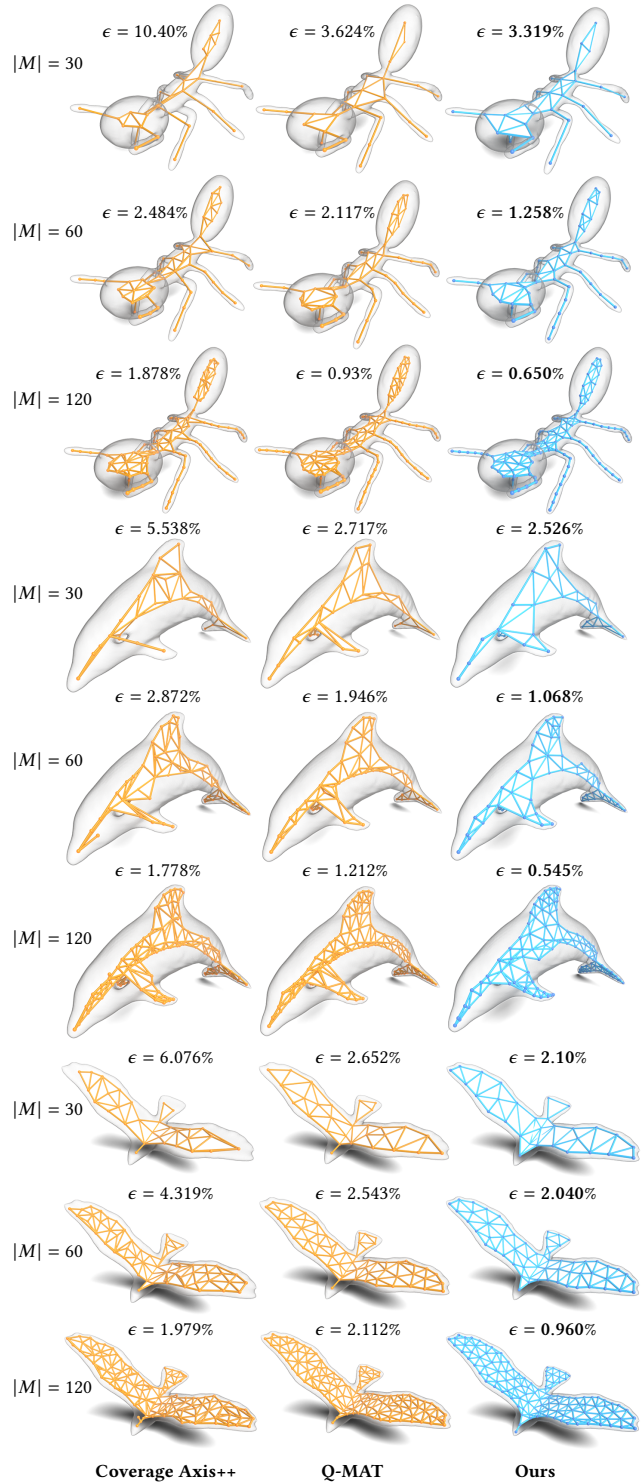


Fig. 10. Comparison of the results obtained with Coverage Axis++, Q-MAT and our method at different resolutions (30, 60 and 120 medial spheres) on three different shapes (bug, dolphin, dove).  $\epsilon$  represents the two-sided Hausdorff distance.