Q-MAT: Computing Medial Axis Transform by Quadratic Error Minimization

Pan Li and Bin Wang
School of Software, Tsinghua University
Feng Sun
Department of Mathematics, Pennsylvania State University
Xiaohu Guo
Department of Computer Science, The University of Texas at Dallas
Caiming Zhang
Department of Computer Science and Technology, Shandong University
and
Wenping Wang
Department of Computer Science, The University of Hong Kong

The medial axis transform (MAT) is an important shape representation for shape approximation, shape recognition, and shape retrieval. Despite years of research, there is still a lack of effective methods for efficient, robust and accurate computation of the MAT. We present an efficient method, called *Q-MAT*, that uses quadratic error minimization to compute a structurally simple, geometrically accurate, and compact representation of the MAT. We introduce a new error metric for approximation and a new quantitative characterization of unstable branches of the MAT, and integrate them in an extension of the well-known quadric error metric (QEM) framework for mesh decimation. Q-MAT is fast, removes insignificant unstable branches effectively, and produces a simple and accurate piecewise linear approximation of the MAT. The method is thoroughly validated and compared with existing methods for MAT computation.

Categories and Subject Descriptors: I.3.5 [**Computer Graphics**]: Computational Geometry and Object Modeling—*Curve, surface, solid, and object representations*

General Terms: Medial Axis, Simplification, Quadratic Error Metric, Stability Ratio

Additional Key Words and Phrases: Volume Approximation

1. INTRODUCTION

The *medial axis transform* (MAT) is a fundamental shape descriptor and has applications in shape approximation, recognition, and retrieval. The *medial axis* of an object in 3D is a set of points with at least two closest points on the object's boundary. It comprises the centers of the spheres that are called *medial spheres*, contained in the object, and touching the object's boundary at two or more points. The MAT [Blum et al. 1967] stores for every point on medial axis its distance to the boundary of the object, i.e., the radius of the associated medial sphere. Hence, the MAT is the combination of the medial axis and the radius function defined on it. Any closed object has a unique MAT and can be reconstructed from its MAT.

We shall assume that the MAT is represented as a 2D simplicial complex, called a *medial mesh*, which is a non-manifold triangle mesh with its vertices being the centers of a set of medial spheres. Through linear interpolation of the mesh spheres along its edges and triangle faces, the medial mesh gives a piecewise linear approximation to the true MAT. Note that such a geometric representation based on sphere interpolation has been adopted in several previous works for geometric modeling [Sun et al. 2013; 2014] and volume approximation [Faraj et al. 2013]. The most commonly-used method for computing a medial mesh of a 3D object is based on computing the Voronoi diagram of a set of sampled points on the object boundary [Amenta and Bern 1998].

A notorious difficulty with the computation of MAT is its instability: the MAT is sensitive to boundary noise or variations; that is, small perturbations to the shape boundary lead to numerous long and unstable branches of the MAT, in the form of dangling line segments or thin and long triangles sticking to the main part of MAT (see the 2D example in Fig. 1). Such unstable branches will be called *spikes*. The spikes are geometrically insignificant because they contribute little to defining the object geometry.

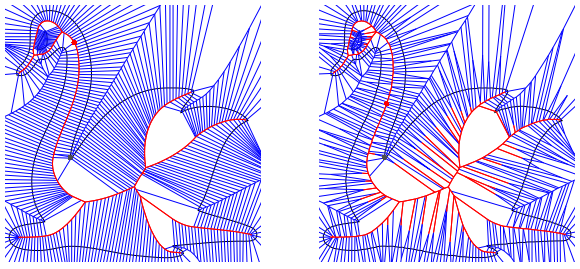

(a) (b)

Fig. 1. The instability of MAT. (a) The medial axis of a swan shape with a smooth boundary (in red). (b) The medial axis of the same shape with slight boundary perturbations. The approximated medial axis (in red) is a subset of the Voronoi diagram (in blue or red).

The instability issue becomes even more severe and difficult to resolve for the MAT of 3D objects. First, spikes may arise due to small variations of the boundary surface. Second, when using the Voronoi-based method to compute the MAT in 3D, spikes may also be generated from sliver tetrahedra in the dual Delaunay triangulation of the Voronoi diagram of the sample points, even if the sample points lie exactly on a noise-free boundary surface [Amenta and

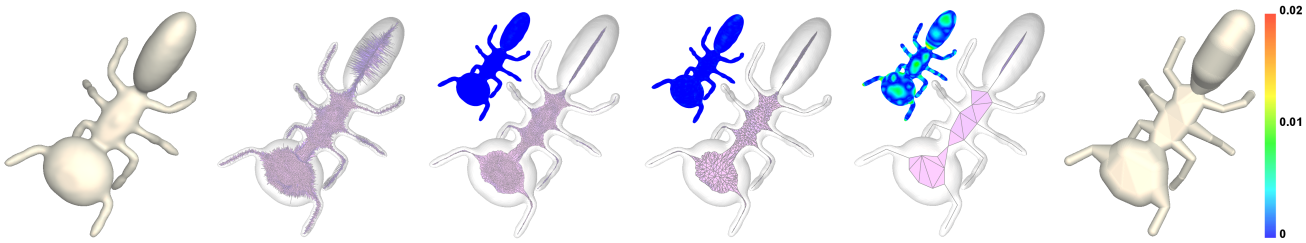

(a) Input shape (b) Initial medial mesh, $\#v = 29,703$, $\epsilon = 0.05\%$ (c) Simplified MAT, $\#v = 13,000$, $\epsilon = 0.28\%$ (d) Simplified MAT, $\#v = 1,000$, $\epsilon = 0.34\%$ (e) Simplified MAT, $\#v = 45$, $\epsilon = 1.98\%$ (f) Reconstructed surface

Fig. 2. Medial axis simplification results by Q-MAT. ($\#v$ is the number of vertices.) (a) The input 3D mesh of Ant model. (b) The initial medial mesh. (c), (d) and (e) The medial meshes at different levels of simplification. (f) The reconstructed surface from the medial mesh in (e). Here, ϵ is the Hausdorff distance from the input surface to the reconstructed surfaces, defined relative to the length of a diagonal of the rectangular bounding box of the input shape. The three color-coded surfaces above the simplified medial meshes show approximation error distributions in Hausdorff distance.

Bern 1998] (see Fig. 6). A sliver is a flat tetrahedron of small volume but has non-degenerate faces. No matter how dense are the boundary sampling points, there exist inevitably many quadruples of adjacent sample points that are nearly co-circular. These quadruple of points define sliver tetrahedra whose centers of circumscribing sphere are far off the true medial axis, thus causing spikes.

For many applications, it is important to remove spikes to yield a structurally simple MAT. Many methods have been proposed for pruning spikes, but they pay little attention on the geometric simplicity and accuracy of the resulting MAT. The requirement on preserving geometric accuracy during MAT simplification is hard to meet because it is difficult to categorically decide which branches are spikes and so should be removed and which are stable so should be kept as geometric features.

Another major issue is the redundancy of the MAT representation. An initial medial mesh computed using a Voronoi diagram-based method is typically a dense and non-manifold mesh – it is structurally complex with many spikes and geometrically redundant with many vertices. While there exist methods for simplifying the MAT by pruning the spikes, it is common to approximate the input 3D shape by the union of medial spheres at the vertices of the medial mesh. That means that a large number of mesh vertices have to be maintained to attain a high approximation accuracy, resulting in a redundant MAT representation. Hence, it is imperative to study how to compute a compact representation of MAT of a small data size. A compact MAT representation is obviously useful for shape compression and helps improve the efficiency of shape deformation, dynamic simulation and shape matching where the MAT is employed.

To summarize, there are three main requirements for a desirable approximate MAT: (1) *structural simplicity*; (2) *geometric accuracy*; and (3) *compactness of representation*. Therefore, the goal of MAT simplification is to both prune spikes and reduce the number of mesh vertices while ensuring the approximation accuracy. In this regard, the existing methods are quite successful in pruning spikes but they tend to remove apparent spikes too aggressively at the expenses of the approximation accuracy. Thus, in general, they perform poorly in preserving approximation accuracy or do not provide a compact MAT representation.

We shall present an efficient method for MAT simplification that meets the above requirements on a high quality MAT approximation. This method, to be called *Q-MAT*, is an extension to the Quadric Error Metric (QEM) framework for mesh decimation by Garland and Heckbert [Garland and Heckbert 1997] and inspired by the recent work on spherical QEM [Thiery et al. 2013]. We adopt the QEM framework due to its simplicity, elegance and superior efficiency. However, several major extensions and modifications are needed to adapt the QEM framework to solving the MAT simplification problem. Unlike an ordinary 2D surface mesh in 3D Euclidean space, the medial mesh is a 2D mesh embedded in 4D space, since every mesh vertex is represented by four coordinates $(x, y, z, r)^\top$, which contain the center $(x, y, z)^\top$ of a medial sphere and its radius r. For shape approximation purposes, the approximate metric on the medial mesh is the Minkowski distance between two spheres, rather than the ordinary Euclidean distance. Hence, we shall reformulate the quadratic error term in the QEM framework using a new metric to properly evaluate and minimize the approximation error for MAT simplification. Furthermore, we use the piecewise linear interpolation of adjacent medial spheres of the medial mesh to approximate the input 3D object. This makes it possible to produce a much simpler medial mesh than merely using the union of medial spheres for approximation, while meeting the same accuracy requirement.

Not all the vertices of a medial mesh are of the same nature. Those vertices involved in defining spikes should be identified and removed in the early stage of simplification to yield a structurally simple and highly accurate simplified medial mesh, while other redundant vertices in a smooth region can be removed at a later stage of simplification if a more compact medial mesh is desired. To achieve this, we introduce a novel measurement, called *stability ratio*, of the geometric significance of each edge of a medal mesh and apply it to prioritizing mesh vertices for simplification in a graded manner, with spike-related vertices removing first.

In summary, the following contributions are made in this paper.

—A new quadratic error metric, called *slab quadratic error metric*, is proposed for measuring approximation errors in MAT simplification;

—A new measure, called *stability ratio*, is proposed for measuring the geometric significance of the edges of a medial mesh. It is used for ranking all the mesh edges so that edges associated with spikes are more likely to be collapsed first.

—A complete method for efficiently simplifying an initial medial mesh to obtain a structurally simple, geometrically accurate, and compact MAT representation.

Fig. 2 shows different simplification levels of the MAT of the Ant model and their reconstruction results. More test examples are presented in a later section for validation.

The remainder of the present paper is organized as follows. Related works on MAT computation are discussed in Section 2. Preliminaries and preparations are given in Section 3 for developing the main algorithm, including the introduction of the stability ratio and the slab quadratic error metric. Our method, Q-MAT, for MAT simplification is presented in Section 4. The experimental results for validation and comparison are presented and discussed in Section 5, and the conclusion is given in Section 6.

2. RELATED WORK

2.1 Medial Axis Simplification Methods

In the following we shall discuss some main existing approaches to MAT computation.

Angle-based filtering method [Attali and Montanvert 1996; Amenta et al. 2001; Foskey et al. 2003; Dey and Zhao 2004; Sud et al. 2005] computes, for every point of the medial axis, the angle formed by its two closest points on the shape boundary. It removes a medial point if the angle associated with that point is less than a user-specified threshold. Angle-based filtering methods usually produce a simplified medial axis with a different topology from the input, although it can preserve local features quite well.

λ-medial axis method [Chazal and Lieutier 2005] uses the circumradius of the closest points of a medial point as a pruning criterion. A medial point is removed if the circumradius associated by the medial point is smaller than a given threshold λ. A small value of λ will preserve the topology of the medial axis. But this criterion cannot preserve features at different scales [Pizer et al. 2003; Siddiqi and Pizer 2008; Attali et al. 2009].

Scale Axis Transform (SAT) [Miklos et al. 2010] prunes spikes more effectively than the previous methods. In SAT, all medial spheres are first scaled by a factor $s > 1$. After scaling, a medial sphere will be removed if it is contained in another medial sphere. Because narrow gaps or small holes may be filled in this step, the homotopy type of the input object may not be preserved by SAT. The final approximate object is obtained by scaling back all the surviving medial spheres by the factor $1/s$. Another shortcoming of SAT is that it can only control the level of simplification by setting the value of s, which is not an intuitive parameter related to the number of mesh vertices to retain. If s is set to be too large, a compact MAT is obtained at the cost of losing many geometric features.

MAT simplification via edge-collapse: The progressive MAT (PMAT) method [Faraj et al. 2013] performs MAT simplification using edge collapse. For an edge connecting two medial spheres, it uses the ratio of the edge length to the difference of the medial radii at the two endpoint as the cost for collapse. However, this method does not optimize the position of medial sphere after each edge-collapse and the results are similar to SAT with little improvement.

An alternative for MAT simplification is the **Hausdorff error-based Method** [Sun et al. 2013; 2014], which defines the error metric by computing the one-sided Hausdorff distance from the original shape to the approximate volumes during simplification. The advantage of the method is its guarantee on approximation accuracy due to its use of the one-sided Hausdorff distance. However, the method has to perform a local exhaustive search for evaluating maximum errors in order to compute the Hausdorff distance, which renders the method slow. Furthermore, because of the difficulty in minimizing an error metric based on the one-sided Hausdorff distance, the vertex position after edge collapse is not placed at an optimal position – one of the endpoints of the collapsed edge is chosen to be the position of the merged point.

Most existing methods for medial axis simplification approximate the medial axis by the union of medial spheres and simplify the MAT using a local or global threshold. In contrast, our Q-MAT method uses the interpolation of medial spheres to define a piecewise approximation of the MAT. Furthermore, the adoption of a novel quadratic error measurement in Q-MAT allows us to achieve fast and accurate MAT simplification.

2.2 QEM and Spherical QEM

Quadric Error Metric Given an input mesh to be simplified, QEM [Garland and Heckbert 1997] uses a quadratic error metric that is the sum of the squared distances from a vertex of a simplified mesh to the containing planes of its associated boundary triangles in the input mesh. The algorithm simplifies the input mesh by evaluating and minimizing these errors between the original input shape and its simplified mesh. Its algorithmic framework takes advantage of the quadratic form of the error terms to allow efficient maintenance and optimization, yielding fast computation.

Spherical Quadric Error Metric (Spherical QEM) [Thiery et al. 2013] extends the QEM framework to extreme simplification of volume representation by computing the squared distance from a sphere to the containing planes of its associated boundary triangles. The spherical QEM method starts by placing the initial spheres (of zero radius) at the vertices of the input surface mesh, and iteratively merges the spheres until a coarse volume approximation of the input shape is produced. Because the Spherical QEM method cannot, in general, produce a valid volume representation at the early stage when the surface approximation error is small, it is not suitable for MAT simplification or high-accuracy volume simplification.

Our method (Q-MAT) is inspired by the QEM method and the Spherical QEM method, but devised for MAT simplification. Unlike the Spherical QEM method that is only for extreme volume simplification, Q-MAT is for MAT simplification, while it is also capable of generating simplified volume representations at different levels of accuracy, including extreme volume simplification.

3. MEDIAL AXIS SIMPLIFICATION

3.1 Medial Mesh

Several previous works use the interpolation of spheres for shape representation. ZSpheres, from ZBrush tool [PIXOLOGIC 2001], uses a manually constructed skeleton of spheres to represent a coarse shape. B-Mesh [Ji et al. 2010] optimizes such a representation for better mesh extraction. Spherical QEM [Thiery et al. 2013], PMAT [Faraj et al. 2013], and the method in [Sun et al. 2013; 2014] all use spheres and their interpolation for volume representation or MAT approximation. We are going to briefly introduce this representation below.

We use a triangle mesh M_s, called *medial mesh* to approximate the MAT of a 3D shape S. Each vertex $\boldsymbol{m}_i$ of M_s represents a medial sphere and is denoted as a 4D point $\boldsymbol{m} = (\boldsymbol{c}, r)^\top$, where $\boldsymbol{c} \in R^3$ is the center of the medial sphere and r its radius. An edge of M_s incident to two medial spheres $\boldsymbol{m}_i$ and $\boldsymbol{m}_j$ is denoted $e_{ij} = \{\boldsymbol{m}_i, \boldsymbol{m}_j\}$. Similarly, a triangle face of M_s is denoted $f_{ijk} = \{\boldsymbol{m}_i, \boldsymbol{m}_j, \boldsymbol{m}_k\}$.

Each edge or face of the medial mesh M_s defines a simple composite volume primitive, as shown in Fig. 3. The primitive given by the edge $\{\boldsymbol{m}_i, \boldsymbol{m}_j\}$ is swept by the family of spheres defined by the linear interpolation of the medial spheres $\boldsymbol{m}_i$ and $\boldsymbol{m}_j$, that is, $(1-t)\boldsymbol{m}_i + t\boldsymbol{m}_j, t \in [0, 1]$. It comprises two spherical caps joined by a truncated cone, and will be called a *medial cone*. The

primitive given by the face $\{\boldsymbol{m}_i, \boldsymbol{m}_j, \boldsymbol{m}_k\}$ is obtained by linearly interpolating the three medial spheres $\boldsymbol{m}_i$, $\boldsymbol{m}_j$, and $\boldsymbol{m}_k$, that is, $a_1\boldsymbol{m}_i + a_2\boldsymbol{m}_j + a_3\boldsymbol{m}_k$, where $a_i (i = 1, 2, 3)$ are the barycentric coordinates with $a_i \geq 0$ and $a_1 + a_2 + a_3 = 1$. This primitive is called a *medial slab* (or just *slab*), bounded by three spherical caps, three conical patches, and two triangles, as shown in Fig. 3 (b).

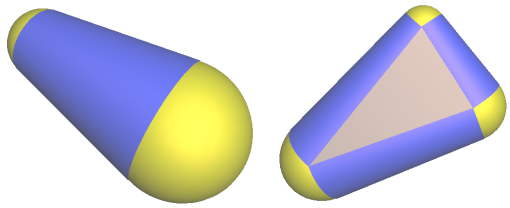

(a) (b)

Fig. 3. Interpolation of spheres. (a) *A medial cone*: the interpolation of two spheres along an edge of the medial mesh. (b) *A medial slab*: the interpolation of three spheres over a triangle face of the medial mesh.

3.2 Slab Quadratic Error

To develop an edge-collapse strategy for MAT simplification, we shall introduce a new error function, called the *slab quadratic error* (SQE), that measures the approximation from the simplified MAT to the initial MAT. SQE is evaluated for each medial sphere of the simplified medial mesh to measure its "distance" to a medial slab of the initial medial mesh in the least squares sense. When performing edge collapse, a new medial sphere results from merging an edge of the medial mesh. The position of this merged sphere is determined by minimizing an error function that is the sum of the SQE terms associated with the two medial spheres to be merged. Similarly to QEM, throughout this iterative edge-collapse process, the SQE function always encodes the approximation error from a medial sphere in the simplified MAT to the initial MAT.

In the following we shall first define the error formula of $SQE_s(\boldsymbol{m}_x)$, for measuring the distance from a medial sphere $\boldsymbol{m}_x$ to a slab s defined by a triangle face of the medial mesh. There are the following requirements on this error measurement. First, the SQE function should be a positive definite quadratic function in the variables, which are the center and radius of the sphere $\boldsymbol{m}_x$, in order to allow efficient minimization in the QEM framework. Second, we require that the function $SQE_s(\boldsymbol{m}_x)$ be minimized if and only if the sphere $\boldsymbol{m}_x$ is tangent to the two planes containing the two boundary triangles of the slab s; that is, $\boldsymbol{m}_x$ is sandwiched between the two bounding planes of the extended slab. This is similar to the error formulation in the QEM method–the error function in QEM measures the squared distance from a point to the plane containing a face triangle, rather than to the triangle itself.

Let an extended slab be defined by its two bounding planes. See Fig. 4. Clearly, the two planes are both tangent to some common sphere $\boldsymbol{m} = (\boldsymbol{c}, r)^\top$ (in red) and denoted by $s = \{\boldsymbol{m}, \boldsymbol{n}_1, \boldsymbol{n}_2\}$, where $\boldsymbol{n}_1, \boldsymbol{n}_2$ are the outward unit normal vectors of the two planes. For the variable sphere $\boldsymbol{m}_x = (\boldsymbol{c}_x, r_x)^\top$ (in blue), we use e_1 and e_2 to measure its distances to the two bounding planes of the slab s. It follows that the squared distances from the variable sphere $\boldsymbol{m}_x = (\boldsymbol{c}_x, r_x)^\top$ (in blue) to the two bounding planes of the slab are

$$\begin{aligned} e_1^2 &= [\boldsymbol{n}_1{}^\top \cdot (\boldsymbol{c} - \boldsymbol{c}_x) + (r - r_x)]^2, \\ e_2^2 &= [\boldsymbol{n}_2{}^\top \cdot (\boldsymbol{c} - \boldsymbol{c}_x) + (r - r_x)]^2. \end{aligned}$$

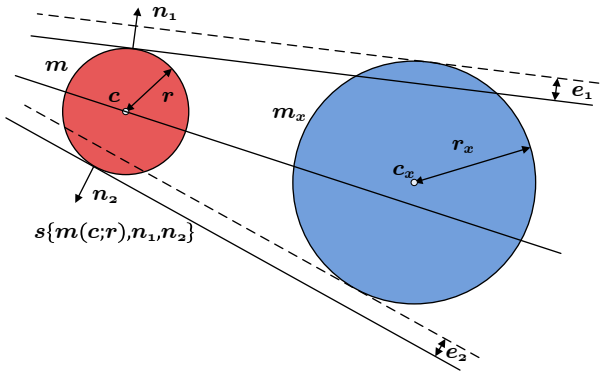

Fig. 4. Illustration for the definition of SQE function. The slab s is specified by a sphere $\boldsymbol{m} = (\boldsymbol{c}, r)^\top$ (in red) and the outward unit normal vectors $\boldsymbol{n}_1$ and $\boldsymbol{n}_2$ of the two bounding planes tangent to $\boldsymbol{m}$. The SQE function $SQE_s(\boldsymbol{m}_x)$ measures the distance from the sphere $\boldsymbol{m}_x = (\boldsymbol{c}_x, r_x)^\top$ (in blue) to the slab s.

Then we define

$$SQE_s(\boldsymbol{m}_x) = e_1^2 + e_2^2. \tag{1}$$

The fact that e_1 and e_1 are linear functions is of great facility, since this enables us to present $SQE_s(\boldsymbol{m}_x)$ as a quadratic function in order to adopt the QEM framework.

To put the above notation in 4D homogenous coordinates, we denote $\overline{\boldsymbol{n}}_1 = (\boldsymbol{n}_1, 1)^\top$, $\overline{\boldsymbol{n}}_2 = (\boldsymbol{n}_2, 1)^\top$, yielding

$$\begin{aligned} SQE_s(\boldsymbol{m}_x) &= [\overline{\boldsymbol{n}}_1^\top \cdot (\boldsymbol{m} - \boldsymbol{m}_x)]^2 + [\overline{\boldsymbol{n}}_2^\top \cdot (\boldsymbol{m} - \boldsymbol{m}_x)]^2 \\ &= \boldsymbol{m}_x{}^\top \cdot \boldsymbol{A} \cdot \boldsymbol{m}_x + \boldsymbol{b}^\top \cdot \boldsymbol{m}_x + c, \end{aligned}$$

where

$$\begin{aligned} \boldsymbol{A} &= \overline{\boldsymbol{n}}_1 \cdot \overline{\boldsymbol{n}}_1^\top + \overline{\boldsymbol{n}}_2 \cdot \overline{\boldsymbol{n}}_2^\top, \\ \boldsymbol{b} &= -2\boldsymbol{A} \cdot \boldsymbol{m}, \\ c &= \boldsymbol{m}^\top \cdot \boldsymbol{A} \cdot \boldsymbol{m}. \end{aligned}$$

Because the MAT of a 3D object can become a curve, such as the MAT of a canal surface, which is approximated by a sequence of edges in a medial mesh, we also need to define a similar quadratic error term for measuring the distance from a sphere to a medial cone, the volume primitive induced by an edge of the medial mesh.

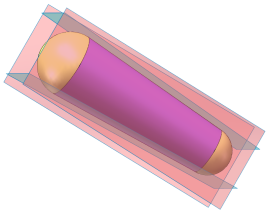

To this end, we bound a medial cone by four planes that are tangent to the cone and symmetric under a rotation of $\pi/2$ about the central axis of the cone, as shown in the right. Since each pair of opposite planes in the four planes forms an extended slab, we sum up the SQE terms of the two resulting slabs to define the error term for the medial cone. Clearly, due to rotational symmetry, although these two bounding slabs of the medial cone are not unique, the final SQE function defined above is independent of the choice of these slabs. This SQE function for a medial cone is also a positive-definite quadratic function and becomes zero if and only if the sphere $\boldsymbol{m}_x$ is inscribed in the medial cone.

We point out that the SQE function is an extension to the *spherical quadric error* proposed in [Thiery et al. 2013] for volume simplification. Over there, the error term measures the distance from

a sphere to a single oriented plane extended by a boundary triangle. Because of the lack of constraint in that formulation, a user-specified radius bound has to be introduced during error minimization to prevent the fitting sphere from going unbounded. In contrast, in our formulation such an unbounded sphere sticking out the slab would lead to a larger SQE value and is therefore prevented by the error-minimization mechanism.

3.3 Stability Ratio

The *spike* is an intuitive, over-simplifying term for describing unstable and geometric insignificant branches of an MAT. To allow a quantitative analysis, we need a measure of this instability to be able to speak about to what degree an edge is stable. Such a measure is also demanded for effectively pruning spikes. In order to achieve progressive MAT simplification to have a structurally simple and accurate medial mesh, it is desirable to first prune spikes as much as possible before decimating the other stable mesh vertices in smooth manifold regions. However, when mesh edges are selected for merging only based on their approximation errors, regardless of whether they are spikes, many stable edges (i.e. non-spike edges) in the dense and smooth regions of the MAT would be selected for collapse equally likely as spikes. Consequently, the spikes would not effectively be pruned until a much later stage of simplification when the approximation error has already become quite large. To address this issues, we shall introduce a quantitative measurement, called *stability ratio*, for ranking mesh edges so that spike edges are prioritized over other edges to collapse first.

Consider an edge $e_{ij} = (\boldsymbol{m}_i, \boldsymbol{m}_j)$ of the medial mesh, where $\boldsymbol{m}_i = (\boldsymbol{c}_i, r_i)$ and $\boldsymbol{m}_j = (\boldsymbol{c}_j, r_j)$ are two spheres being the endpoints of e_{ij}. The distance between the two spheres is [Choi and Seidel 2001]

$$d_h(\boldsymbol{m}_i, \boldsymbol{m}_j) = max\{0, ||\boldsymbol{c}_i - \boldsymbol{c}_j|| - |r_i - r_j|\}. \quad (2)$$

Note that $||\boldsymbol{c}_i - \boldsymbol{c}_j|| - |r_i - r_j| < 0$ if and only if one sphere is contained inside the other. The *stability ratio* of the edge e_{ij} is then defined as the ratio of this distance $d_h(\boldsymbol{m}_i, \boldsymbol{m}_j)$ to the Euclidean distance between the centers of the two spheres, that is,

$$\Gamma_{ij} = \frac{d_h(\boldsymbol{m}_i, \boldsymbol{m}_j)}{||\boldsymbol{c}_i - \boldsymbol{c}_j||}. \quad (3)$$

Fig. 5 shows $d_h(\boldsymbol{m}_1, \boldsymbol{m}_2)$ as the difference between the two spheres $\boldsymbol{m}_1$ and $\boldsymbol{m}_2$. Clearly, Γ_{ij} gives a continuous measurement on the stability of an edge, with $0 \leq \Gamma_{ij} \leq 1$. When Γ_{ij} is 0 for two distinct spheres, the smaller sphere is internally tangent to the larger one and the edge e_{ij} is a *pure* spike; when $\Gamma_{ij} = 1$, the radii of the two spheres are equal and the edge e_{ij} is a completely stable edge, that is, not a spike. In Fig. 6 we use false color to display the stability ratios of all the edges of two unprocessed medial meshes computed using the Voronoi-based method.

4. SIMPLIFICATION ALGORITHM

4.1 Outline of Algorithm

Our MAT simplification algorithm (Q-MAT) follows the QEM framework [Garland and Heckbert 1997] based on edge-collapse. All the edges are queued according to their cost values, and the edge with minimum cost is iteratively collapsed to generate a new medial sphere $\boldsymbol{m}_g$ as the result of merging the two endpoints of the edge. The flow of the algorithm is shown as Algorithm 1.

In the following subsections, we shall describe how to optimize the position of the merged medial sphere $\boldsymbol{m}_g$ when an edge is collapsed and how to compute the collapse cost c_{ij} of an edge e_{ij} to determine its priority to be merged. Note that all the coefficients and parameter values to appear in the expressions below are given for models that are scaled by setting the diagonal length of their bounding boxes to be 1.

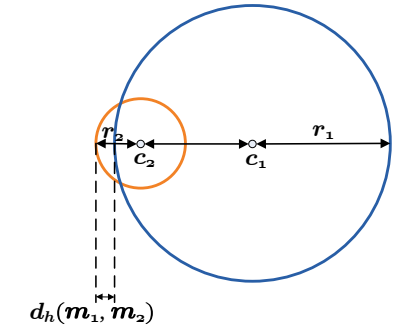

Fig. 5. Illustration of the distance of two spheres.

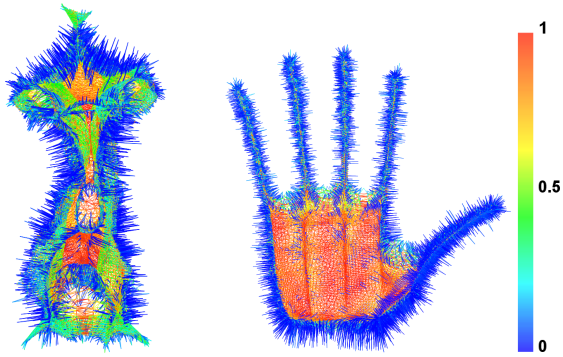

Fig. 6. False-color display of the stability ratios of the edges for two medial meshes. The blue color (towards 0) indicates unstable edges, i.e. spike edges, and the red color (towards 1) indicates stable edges.

4.2 Positioning a New Medial Sphere for Edge Collapse

At initialization, each medial sphere $\boldsymbol{m}$ is assigned an error function $E_{\boldsymbol{m}}(\boldsymbol{m}_x)$ which is the sum of the SQE terms of $\boldsymbol{m}$ with respect to the set of all the slabs that $\boldsymbol{m}$ is incident to, denoted by the set $slabs(\boldsymbol{m})$. This error function $E_{\boldsymbol{m}}(\boldsymbol{m}_x)$ can be written as

$$E_{\boldsymbol{m}}(\boldsymbol{m}_x) = \sum_{s \in slabs(\boldsymbol{m})} SQE_s(\boldsymbol{m}_x). \quad (4)$$

Each edge $e_{ij} = \{\boldsymbol{m}_i, \boldsymbol{m}_j\}$ is assigned a cost function $E^e_{\boldsymbol{m}_i, \boldsymbol{m}_j}(\boldsymbol{m}_x)$ defined as the sum of the error functions associated with the two medial spheres at its endpoints:

$$E^e_{\boldsymbol{m}_i, \boldsymbol{m}_j}(\boldsymbol{m}_x) = E_{\boldsymbol{m}_i}(\boldsymbol{m}_x) + E_{\boldsymbol{m}_j}(\boldsymbol{m}_x). \quad (5)$$

The optimal position of the merged sphere $\boldsymbol{m}_g$ of $\boldsymbol{m}_i$ and $\boldsymbol{m}_j$ is computed as the minimizer of $E^e_{\boldsymbol{m}_i, \boldsymbol{m}_j}(\boldsymbol{m}_x)$, that is,

$$\boldsymbol{m}_g = \arg\min_{\boldsymbol{m}_x} E^e_{\boldsymbol{m}_i, \boldsymbol{m}_j}(\boldsymbol{m}_x). \quad (6)$$

That is, we use the sphere $\boldsymbol{m}_g$ that best approximates the set of the involved slabs of $slabs(\boldsymbol{m}_i) \cup slabs(\boldsymbol{m}_j)$ to replace the edge e_{ij}. The edge collapse cost function of e_{ij} is then set to be

$edge_{cost}(e_{ij}) = E^e_{\boldsymbol{m}_i,\boldsymbol{m}_j}(\boldsymbol{m}_g)$; the edge costs of all the edges are ranked to determine which edge is to collapse first. If the edge e_{ij} is collapsed, the merged sphere $\boldsymbol{m}_g$ becomes associated with all the slabs in the union of $slabs(\boldsymbol{m}_i)$ and $slabs(\boldsymbol{m}_j)$, that is,

$$slabs(\boldsymbol{m}_g) = slabs(\boldsymbol{m}_i) \cup slabs(\boldsymbol{m}_j). \quad (7)$$

Note that the set $slabs(\boldsymbol{m}_g)$ of the slab is not recorded in the algorithm. The error function $E_{\boldsymbol{m}_g}(\boldsymbol{m}_x)$ of $\boldsymbol{m}_g$ that properly measures the approximation error from a variable sphere $\boldsymbol{m}_x$ to $slabs(\boldsymbol{m}_g)$ is obtained by setting

$$E_{\boldsymbol{m}_g}(\boldsymbol{m}_x) = E^e_{\boldsymbol{m}_i,\boldsymbol{m}_j}(\boldsymbol{m}_x).$$

When minimizing the quadratic function

$$E^e_{\boldsymbol{m}_i,\boldsymbol{m}_j}(\boldsymbol{m}_x) = {\boldsymbol{m}_x}^\top \cdot \boldsymbol{A} \cdot \boldsymbol{m}_x + \boldsymbol{b}^\top \cdot \boldsymbol{m}_x + c,$$

we consider the following cases:

—If $\boldsymbol{A}$ is invertible, we compute the optimal sphere as $\boldsymbol{m}_g = -\frac{1}{2} \cdot \boldsymbol{A}^{-1} \cdot \boldsymbol{b}$;

—If $\boldsymbol{A}$ is not invertible, one way is to find an optimal sphere along the edge as done similarly in [Thiery et al. 2013]. However, for simplicity, we just select the sphere with minimum $E^e_{\boldsymbol{m}_i,\boldsymbol{m}_j}(\boldsymbol{m}_x)$ from the three spheres: $\boldsymbol{m}_i, \boldsymbol{m}_j$, and $(\boldsymbol{m}_i + \boldsymbol{m}_j)/2$.

In addition, we apply some additional treatments to prevent the shrinking of the boundary edges of the medial mesh, as we will describe in Section 4.4.

4.3 Modified Edge Collapse Cost

As discussed in Section 3.3, if the mesh edges are selected for collapsing based on their approximation errors only, some spikes will not be pruned until a much later stage when the approximation error is already quite large. Consequently, we may then get a structurally simple medial mesh but it is not an accurate approximation any more. Recall that we have introduced the *stability ratio* to measure the stability of a medial edge. In order to prioritize unstable edges to collapse first, we select the edges to collapse based on a modified collapse cost, which is a combination of the SQE function and the *stability ratio*. This new cost function is defined as

$$c_{ij} = (E^e_{\boldsymbol{m}_i,\boldsymbol{m}_j}(\boldsymbol{m}_g) + k) * {\Gamma_{ij}}^2, \quad (8)$$

where $E^e_{\boldsymbol{m}_i,\boldsymbol{m}_j}(\boldsymbol{m}_g)$ is the original edge cost defined previously in the last subsection, k a positive constant, and Γ_{ij} the *stability ratio*.

The rationale behind the formulation of this new collapse cost is as follows. At the initial stage, since the approximation error is nearly zero and therefore much smaller than the constant k, this new cost formulae makes us select the edges to collapse mainly according to the values of stability ratio Γ_{ij}; that is, those edges with smaller stability ratios will be selected first. Hence, most spikes will get pruned in the early stage. As the simplification progresses, the approximation error $E^e_{\boldsymbol{m}_i,\boldsymbol{m}_j}(\boldsymbol{m}_g)$ will be larger due to error accumulation and eventually become the dominating term of the cost function c_{ij}. By that time, the edges will be selected mainly according to the values $E^e_{\boldsymbol{m}_i,\boldsymbol{m}_j}(\boldsymbol{m}_g)$, since most spikes, which have relatively small stability ratio Γ_{ij}, have already been removed and the stability ratios of the remaining edges are relatively large. Hence, in this stage, the stable edges with smaller approximation errors, especially those interior edges in the dense and smooth regions of the mesh, will be selected to be collapsed.

Algorithm 1: Medial Axis Simplification

```
Input: M_s = {{m_i}, {e_ij}}, the medial axis of S
Input: σ, the number of medial spheres; when reaching this,
       the algorithm ends
Input: η, the number of medial spheres; when reaching this,
       the topology check starts
Output: M̄_s = {{m_i}, {e_ij}}, simplified medial axis of S
// Q is a priority queue of edges sorted by
   cost;
// n_o is the number of the original medial
   spheres;
// n_r is the remaining number of medial
   spheres;
n_r = n_o;
for each edge e_ij do
    compute the optimal contraction target m_g;
    compute the collapse cost c_ij;
    Q ← e_ij with m_g and c_ij;
end
while Q not empty and n_r > σ do
    e_ij ← Q.top() ;
    Q.pop();
    if m_i is valid and m_j is valid then
        if n_r <= η then
            topology preservation check;
        end
        collapse e_ij → m_g;
        make all neighbor of m_i and m_j adjacent to m_g;
        mark m_i, m_j as invalid;
        for each neighbor m_k of m_g do
            compute the optimal contraction target m_t;
            compute the collapse cost c_kg;
            Q ← e_kg with m_t and c_kg;
        end
        n_r- - ;
    end
end
```

Certainly, the constant k is an important parameter affecting the behavior of this edge selection strategy. We now use the Plane model in Fig. 7 to show the influence of k. It is clear that the stability ratios of the boundary edges are smaller than those of the interior edges, as shown in Fig. 7(a). As a consequence, as shown in Fig. 7, a too large value of k causes to prune the spikes faster, since it delays the time when the error term $E^e_{\boldsymbol{m}_i,\boldsymbol{m}_j}(\boldsymbol{m}_g)$ becomes dominant. However, after pruning the spikes, such a large k makes the simplification focus more on the edges with smaller stability ratio, thus generating a medial mesh with non-uniform vertex distribution, as shown in Fig. 7(b) and (c). We can see from Table I that the mesh quality of Fig. 7(d) is clearly better than that in (b) and (c). When the value of k is too small, the influence of $E^e_{\boldsymbol{m}_i,\boldsymbol{m}_j}(\boldsymbol{m}_g)$ is overemphasized at the beginning, leading to more uniform vertex distribution of the simplified medial axis but at the cost of delaying the pruning of spikes, as shown in the Fig. 8(d). Based on our extensive testing, including those cases shown in Fig. 7, we choose $k = 10^{-5}$ in the implementation of Q-MAT for generating all the results in Section 5.

We next use the Bear model in Fig. 8 as an example to show the necessity of incorporating k and the stability ratio Γ_{ij} in the new cost function in Eq. 8. Fig. 8 shows the simplification processes

Table I. Quality Comparison of Simplified Medial Mesh

Model	$\bar{\theta}_{min}$	$\bar{\theta}_{max}$	$\theta < 30^\circ$
Fig.7(b)	18.55	108.98	78.3%
Fig.7(c)	20.99	106.55	72.2%
Fig.7(d)	25.19	101.76	61.1%

$\bar{\theta}_{min}$ is the average of minimal angles of each triangle in the medial mesh; $\bar{\theta}_{max}$ is the average of maximal angles of each triangle in the medial mesh; $\theta < 30^\circ$ is the percentage of triangles with its minimal angle smaller than 30 degrees.

without using k (i.e. setting $k = 0$) or without using k and Γ_{ij} (i.e. setting $\Gamma_{ij} = 1$), respectively.

Figures 8(b) and (c) show that if we do not use k then the pruning of spikes will be delayed and we cannot get a smooth medial mesh until simplifying the medial mesh down to 3000 vertices. Figures 8(d), (e) and (f) show the simplification results using only the error term $E^e_{\boldsymbol{m}_i,\boldsymbol{m}_j}(\boldsymbol{m}_g)$ (i.e. without k and Γ_{ij}) in Eq. 8. In this case we cannot get a smooth medial mesh even when the MAT is simplified to have only 500 vertices. In comparison, Fig. 17 shows the simplification process using Eq. 8, with $k = 10^{-5}$. In this case the spikes are removed at the early stage while maintaining good mesh quality throughout. Here, a compact medial mesh is already produced when the MAT is simplified to have 10,000 vertices.

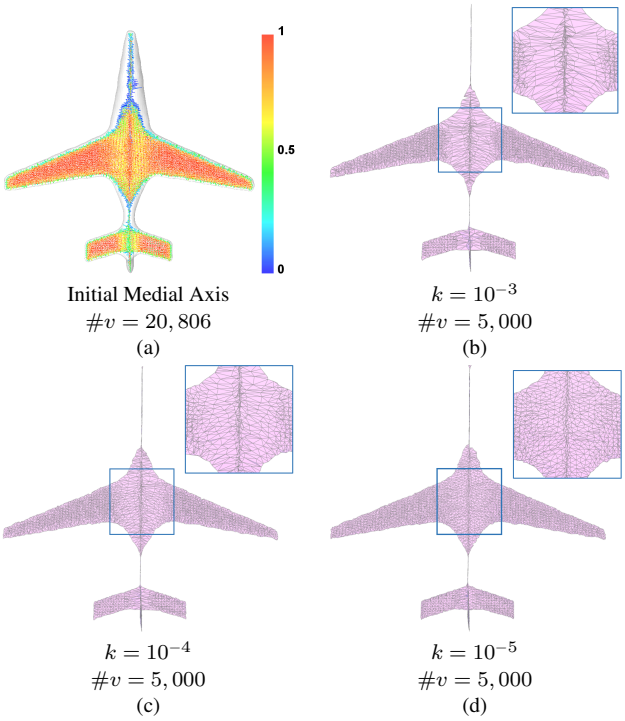

Initial Medial Axis $\#v = 20,806$ (a)
$k = 10^{-3}$ $\#v = 5,000$ (b)
$k = 10^{-4}$ $\#v = 5,000$ (c)
$k = 10^{-5}$ $\#v = 5,000$ (d)

Fig. 7. Influence of the importance parameter k. (a) The initial medial mesh with edges displayed using false-color of the stability ratios. (b), (c) and (d): The simplified medial meshes with $5,000$ vertices generated using different values of k.

4.4 Boundary Preservation

The *boundary edges* of the medial axis are the edges incident to only one face or no face in the medial mesh. The MAT in general consists of 2D non-manifold surfaces and 1-D curves and therefore has open boundaries. A problem with QEM-like simplification is

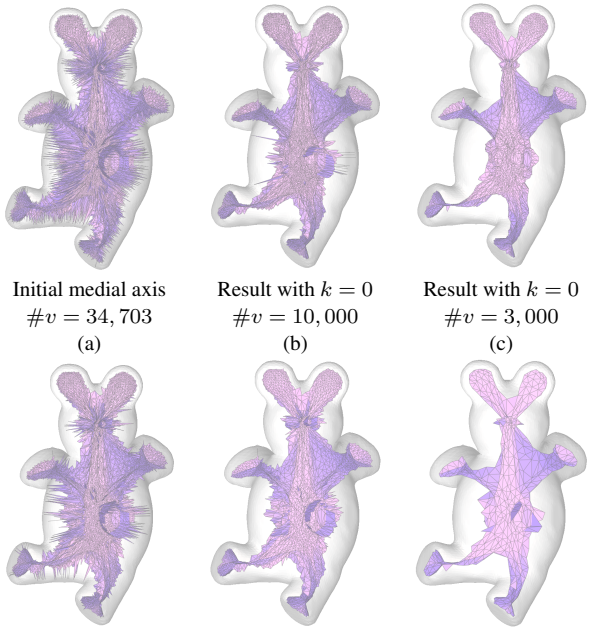

Initial medial axis $\#v = 34,703$ (a)
Result with $k = 0$ $\#v = 10,000$ (b)
Result with $k = 0$ $\#v = 3,000$ (c)
With only $E^e_{\boldsymbol{m}_i,\boldsymbol{m}_j}$ $\#v = 10,000$ (d)
With only $E^e_{\boldsymbol{m}_i,\boldsymbol{m}_j}$ $\#v = 3,000$ (e)
With only $E^e_{\boldsymbol{m}_i,\boldsymbol{m}_j}$ $\#v = 500$ (f)

Fig. 8. Simplification results with different conditions. (a) The initial medial mesh. (b) and (c): The simplified medial mesh by setting $k = 0$ in Eq. 8. (d), (e) and (f): The simplified medial mesh with $\Gamma_{ij} = 1$ and $k = 0$ in Eq. 8;

that some stable boundary edges of MAT may shrink during simplification. Hence, we will introduce additional constraints to prevent the stable boundary edges of the medial mesh from shrinking. There are two kinds of boundary elements to consider: (1) a boundary vertex $\boldsymbol{m}_v$ (which is a boundary medial sphere) is incident to exactly one mesh edge e that is not incident to any triangle face; (2) a boundary edge e (which is a boundary medial cone) incident to exactly one triangle face f.

In the first case, as shown in Fig. 9(a), the *medial cone* defined by the edge e incident to the boundary vertex $\boldsymbol{m}_v$ (shown as the blue sphere) is already bounded by four planes tangent to it. We just need to add an oriented plane p tangent to the boundary medial sphere $\boldsymbol{m}_v$ and orthogonal to the edge e. In the second case, as a slab the face f incident to the boundary edge e is already bounded by two planes. Then we just need to add an oriented plane p tangent to the medial cone defined by the edge e and orthogonal to the triangle f, as shown in Fig. 9(b).

The error terms defined by the squared distances from a sphere to the added planes are then included in the SQE terms of the involved medial spheres as "soft" constraints in the simplification process to help prevent the stable boundary of the MAT from shrinking. Specifically, at initialization, for a boundary medial sphere $\boldsymbol{m}_v$ in the first case above, the corresponding error term is added to the error function associated with $\boldsymbol{m}_v$; for a boundary medial cone defined by an edge e in the second case above, the corresponding error term is added to the error function associated with each of the two medial spheres at the endpoints of the edge e.

Because unstable boundary vertex or edges (i.e. spikes) should not be protected, we weigh these added error terms with the squared *stability ratio* so that those error terms associated with unsta-

ble boundary vertices or edges will have smaller weights. Consequently, unstable boundary vertices or edges will not be protected as much as stable boundary vertices or edges. Specifically, we let the added quadratic error terms associated with boundary vertices or edges take the form

$$\bar{e}^2 = k' * {\Gamma_{ij}}^2 * d_p^2(\boldsymbol{m}_x),$$

where $d_p^2(\boldsymbol{m}_x)$ is the squared distance from the variable sphere $\boldsymbol{m}_x$ to the added bounding plane p. We set $k' = 0.1$ in all our experiments. Hence, these error terms are first weighted before they are added to the SQE functions of the involved medial spheres. We have observed that the stable boundaries of the MAT are well preserved with this scheme, without compromising the effectiveness or efficiency of the overall method.

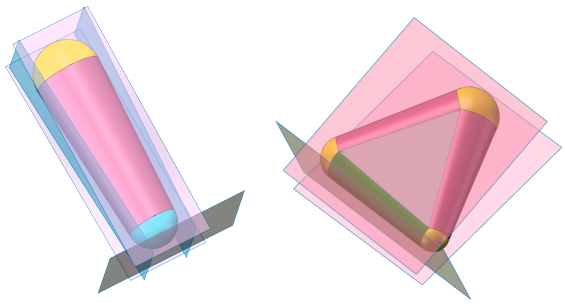

(a) (b)

Fig. 9. Boundary preservation strategy. (a) Protection of a dangling edge of a medial mesh; the blue sphere represents the boundary medial sphere. (b) Protection of a boundary edge of a medial mesh; the green cone represents the boundary medial edge. These added green planes make the merged sphere stay close to the boundary.

4.5 Topology preservation

To preserve the topology of the medial mesh, we use the strategy of hole detection and link condition detection similar to the method of Dey et al. [1998], which discusses edge contractions in simplicial complexes and local conditions for topology preservation, i.e. preventing a hole from disappearing due to edge contraction. We apply this strategy only when the number of vertices of simplified MAT is reduced to 200, since almost all spikes have been removed at that stage and this treatment works well for all of our experiments. With this strategy, the simplification process will stop when further edge contraction will lead to topological change. Taking the Fertility model in Fig. 10 for example, we need at least 9 medial cones to preserve the topology and any further contraction will make a loop of the MAT disappear. Small loops in the MAT are also well preserved with this strategy. As shown in Fig. 11. Here the small loops at the tip of the weapon in the Neptune model are preserved.

As default, we start the topology check when there are 200 vertices left. Of course, this default setting does not work for a model with genus larger than 200, for example. In such cases, the user may specify when to start the topology check. We do not add this topology check step from the very start of simplification because it is quite time consuming–the time needed for topological checking is about 3 to 5 times longer than the current running time.

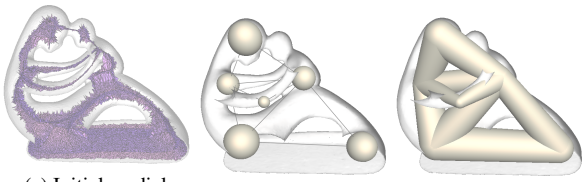

(a) Initial medial mesh $\#v = 153,695$ (b) $\#v = 6$ (c) $\#v = 6$

Fig. 10. Simplification results with topology preservation. (a) The initial medial mesh. (b) The simplified medial mesh. (c) The reconstructed shape from (b).

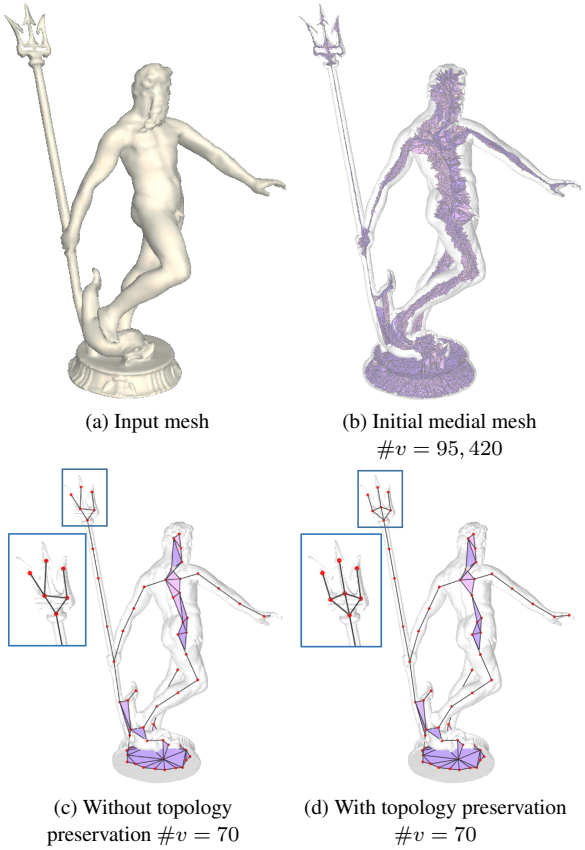

(a) Input mesh (b) Initial medial mesh $\#v = 95,420$

(c) Without topology preservation $\#v = 70$ (d) With topology preservation $\#v = 70$

Fig. 11. Effectiveness of the topology preservation scheme. (a) The input boundary mesh. (b) The initial medial mesh. (c) The simplified result without checking for topology preservation. (d) The simplified result with topology preservation.

4.6 Preventing Mesh Inversion

Collapsing an edge $e_{ij} = \{\boldsymbol{m}_i, \boldsymbol{m}_j\}$ may cause some of its neighboring faces to fold over. A method for avoiding this kind of undesirable mesh inversion has been discussed in the original QEM method [Garland and Heckbert 1997]. When this occurs, we adopt the following strategy. We consider the three options of placing the merged vertex at $\boldsymbol{m}_i$, $\boldsymbol{m}_j$, or $(\boldsymbol{m}_i + \boldsymbol{m}_j)/2$, and choose the one that attains the smallest collapse cost among the three options *and* is free of mesh inversion. If none of the options is free of mesh inversion, the edge e_{ij} will not be collapsed.

5. RESULTS AND DISCUSSIONS

In this section we will show the simplification results of Q-MAT for validation and comparisons with other methods. We implement Q-MAT in C++ and our program runs on a Windows 7 workstation with an Intel i5 CPU @2.50GHz and 8 GB memory. All the models are processed after being scaled such that the diagonal lengths of their bounding boxes are equal to 1.

5.1 Preparation of Initial MAT

In the following we will show the results of using Q-MAT to simplify initial MATs that are output by the SAT method [Miklos et al. 2010] and by the method of Amenta and Bern [1998], respectively. In the former case, we use the publicly available executable program of SAT provided by the authors and use its output as the initial MAT for Q-MAT. SAT incrementally builds a set of sample points and an approximating surface mesh until the approximation tolerance is met, then identifies the inner poles from the constructed Voronoi diagram. As for the output of the method of Amenta and Bern [1998], we make the following adjustment. While the method of Amenta and Bern [1998] provides a way of filtering the Voronoi diagram using the "poles" of the Voronoi cells to improve the stability of the MAT, we take the original Voronoi cells without using any filtering, since we wish to test how Q-MAT can handle highly complex initial MAT with many spikes. The Voronoi diagram is computed by first getting the Delaunay triangulation of the sample points using the CGAL package "Delaunay Triangulation 3" and then taking its dual. In this way, we get the initial medial mesh with the medial spheres connected to its neighbors.

We use the one-sided Hausdorff distance error, denoted ϵ, to assess the approximation accuracy of the simplified medial meshes to be presented in this section. For a simplified medial mesh, this error is defined to be the maximum of the shortest distances from a set of densely sampled points on the input surface to the surface reconstructed from the simplified medial mesh and is expressed in percentages with respect to the diagonal length of the input object' bounding box. The reconstructed shape can be got by rendering all the primitives including medial spheres, medial cones and medial slabs, so the shortest distance from a point to the reconstructed surface can be computed by finding its shortest distance to the surfaces of all the primitives.

5.2 Comparisons and Discussions

We compare our simplification method, Q-MAT, with the following four methods: (1) the angle-based method [Foskey et al. 2003]; (2) the λ-medial axis method [Chazal and Lieutier 2005]; (3) the SAT method [Miklos et al. 2010]; and (4) the method [Sun et al. 2013; 2014], which will be referred to as HEM because it uses the one-sided Hausdorff error metric to control the simplification of a medial mesh. For the SAT and HEM methods, we use the codes provided by the authors. We use our own implementations of the angle-based method and the λ-medial axis method.

Comparison with angle-based simplification and λ-medial axis: Fig. 12 shows the comparison of Q-MAT with the angle-based method and the λ-medial axis method. Starting with an initial MAT of $74,998$ vertices of the Venus model, we use the three methods to simplify it to produce three medial meshes of $20,000$ vertices and three medial meshes of $10,000$ vertices, respectively. For each simplified medial mesh, we also show the reconstructed object and its color-coded distribution of the one-sided Hausdorff error between the input mesh and the reconstructed objects relative to the diameter of the bounding box.

For the simplified medial mesh of $20,000$ vertices produced by the angle-based method (Fig. 12(b)), there still remain unstable branches while the medial mesh has already become disconnected. For the simplified medial mesh of $20,000$ vertices produced by the λ-medial axis method (Fig. 12(d)), there are many holes in the medial mesh, and the holes become larger as the medial mesh is further simplified to have $10,000$ vertices (Fig. 12(e)). In contrast, Q-MAT produces more accurate approximations to the MAT with smaller Hausdorff errors ((Fig. 12(f) and (g)).

Comparison with the SAT method: Here an initial MAT of $140,607$ vertices of a vase model is generated using the same method used for generating initial MATs in the original paper [Miklos et al. 2010] on the SAT method. We experimented with the SAT method using two values of the scaling parameter: $s = 1.1$ and $s = 1.5$, to generate two simplified medial meshes. Then we ran Q-MAT to generate two medial meshes with the same number of vertices, respectively. Fig. 13(c) shows the medial mesh by the SAT method with $s = 1.1$. It has $91,325$ vertices and is structurally cleaner than the simplified result by Q-MAT, which is shown in Fig. 13(d), while the result by Q-MAT is more accurate. The medial mesh by the SAT method with $s = 1.5$ has $38,921$ vertices. It has become structurally incomplete and has a large approximation error, as shown in Fig. 13(e). In contrast, the simplified medial mesh by Q-MAT with the same number of vertices is structurally simpler and still very accurate, as shown in Fig. 13(f). The medial mesh is further simplified by Q-MAT to have $10,000$ vertices and 250 vertices, as shown in Fig. 13(g) and (h), respectively. Note that it is difficult to choose a value of the scaling factor s in the SAT method to produce such extreme simplification of the medial mesh having only 250 vertices and having the matching approximation quality. Hence, Q-MAT is more suitable than the SAT method for progressive MAT simplification.

Comparison with the Hausdorff error-based method (HEM): We use the Octopus model in Fig. 14 to compare Q-MAT with HEM [Sun et al. 2013; 2014]. Starting from the same initial MAT of $39,342$ vertices, we use both simplification methods to reduce the vertex number to $18,000$, $1,000$, 100 and 50, respectively. We will only show the color-coded results for the stage of 100 and 50 vertices because the errors for the first two cases ($18,000$ and $1,000$) are very small. It can be seen that Q-MAT removes almost all the spikes when there are $18,000$ vertices left, while HEM fails to do so even when the number of vertices has been reduced to $1,000$.

Since HEM always chooses to collapses the edge that has the smallest Hausdorff approximation error, the Hausdorff distance is generally smaller than Q-MAT at the beginning. However, at the extreme level of simplification, i.e. when $\#v = 100$ or 50, Q-MAT produces smaller Hausdorff errors than HEM, as observed consistently on many models we have tested. That is mainly because the position of the merged medial sphere after each collapse is more carefully optimized in Q-MAT than in HEM. Moreover, computing the Hausdorff distance to evaluate local maximum errors for makes HEM relatively slow. The total simplification time of HEM for the Octopus model from its initial state to 100 vertices is about $230s$, while Q-MAT takes less than $6s$. Table II gives the simplification time of Q-MAT from the initial medial meshes of various objects to simplified medial meshes of 100 vertices.

5.3 More Results

To test Q-MAT in some special cases with known ground truths, we show an example in Fig. 15 of using Q-MAT to simplify the MAT of a mesh surface approximating a unit sphere centered at the

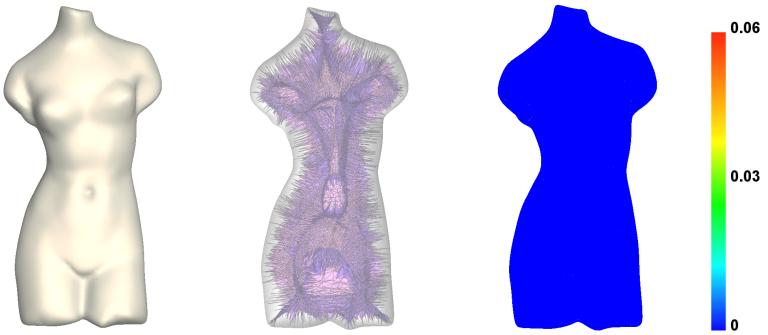

(a) Venus model($\#v = 74,998, \epsilon = 0.05\%$)

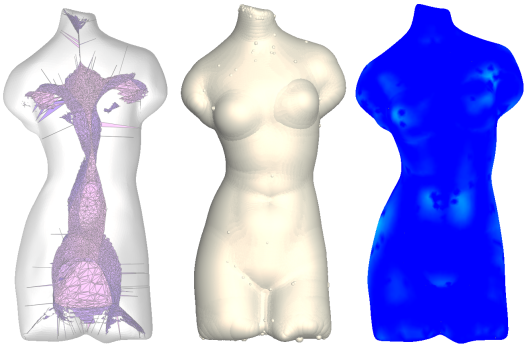

(b) angle-based method($\#v = 20,000, \epsilon = 1.35\%$)

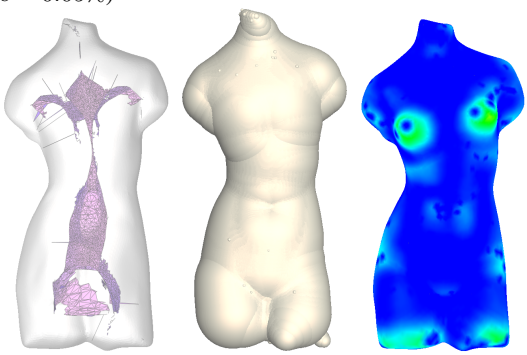

(c) angle-based method($\#v = 10,000, \epsilon = 3.17\%$)

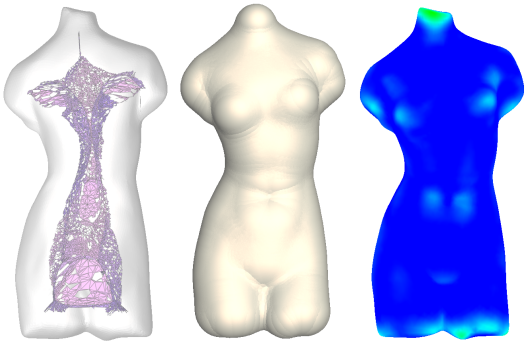

(d) λ-medial axis($\#v = 20,000, \epsilon = 2.52\%$)

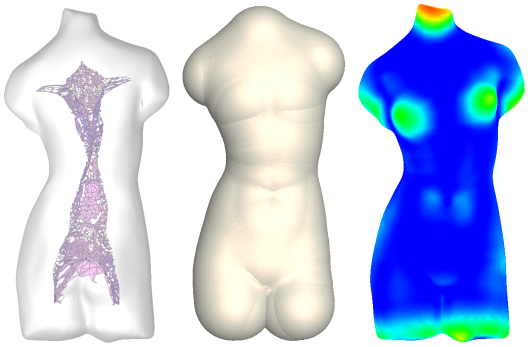

(e) λ-medial axis($\#v = 10,000, \epsilon = 5.52\%$)

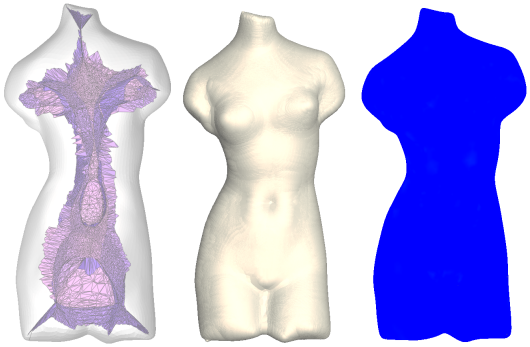

(f) Q-MAT ($\#v = 20,000, \epsilon = 0.25\%$)

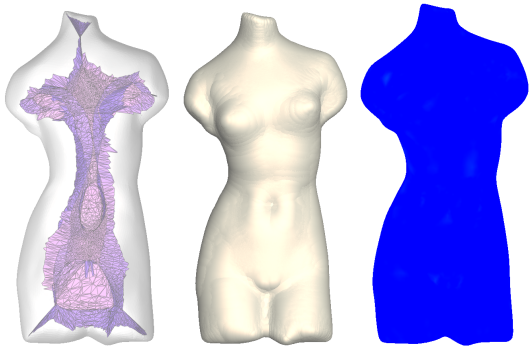

(g) Q-MAT ($\#v = 10,000, \epsilon = 0.33\%$)

Fig. 12. The comparison of Q-MAT with the angle-based method and the λ-medial axis method. ($\#v$ is the number of mesh vertices.) (a) The input mesh of the Venus model, the initial medial mesh, and the color-coded distribution of the errors from the input mesh to the reconstructed surface. (b) and (c): Simplification results by the angle-based method. (d) and (e): Simplification results by the λ-medial axis method. (f) and (g): Simplification results by Q-MAT. In each group, from left to right, there are the simplified medial mesh, the reconstructed surface from it, and the color-coded distribution of the errors from the input surface to the reconstructed surface.

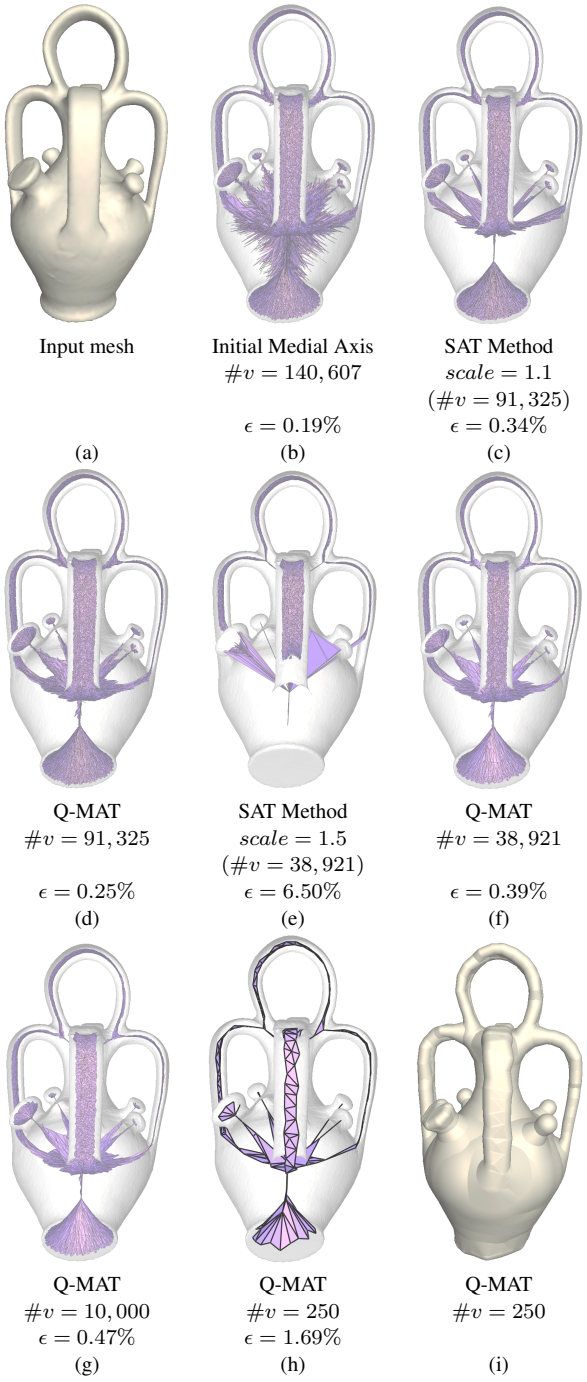

Fig. 13. The comparisons of Q-MAT with the SAT method. ($\#v$ is the number of mesh vertices.) (a) The input mesh. (b) The initial medial mesh. (c) The simplified medial mesh by SAT with $s = 1.1$. (d) The simplified medial mesh by Q-MAT. (e) The simplified medial mesh by SAT with $s = 1.5$. (f) The simplified medial mesh by Q-MAT. (g) and (h): Two further simplified medial meshes by Q-MAT. (i) The shape reconstructed from the medial mesh in (h).

origin, with $2,562$ vertices. The initial medial mesh has $7,156$ vertices. Q-MAT simplifies the initial medial mesh to a single medial sphere centered at $(10^{-9}, -10^{-9}, -4.0 * 10^{-10})$ with radius equals to 0.9999997, which is a good approximation to the unit sphere.

Table II. Comparison of HEM and Q-MAT

Model (#v)	HEM (s)	Q-MAT (s)	Ratio of time
Ant (29,703)	142.3	4.0	35.6
Bear (34,703)	229.7	5.5	41.8
Venus (74,998)	684.0	11.9	57.5
Plane (20,806)	118.2	2.5	47.3
Dolphin (54,241)	386.2	9.5	40.7
Octopus (39,342)	230.0	5.4	42.6

"#v" is the number of vertices in the initial medial mesh. The first column shows the model, the second column shows the simplification time of HEM, the third column shows the simplification time of Q-MAT, and the fourth column shows the ratios of simplification time of HEM over to that of Q-MAT. The simplification time is from the initial MAT to getting the simplified MAT with 100 vertices.

Robustness to noise: Since boundary surface noises tend to cause spikes, we wish to know whether Q-MAT still works effectively in the presence of increased levels of boundary surface noises. In Fig. 16, we add random displacements for each vertex of the Bird model, in different percentages with respect to the diagonal length of the bounding box, denoted by $\eta \in [0, 1]$. Although Q-MAT tends to treat the noises as details of the shape, as simplification progresses, it is capable of distinguishing the important parts of the MAT and preserving them.

Fig. 17, Fig. 18, and Fig. 19 show some more MAT simplification examples by Q-MAT. In the examples of Fig. 19, we use the initial medial mesh from SAT [Miklos et al. 2010]. For other examples, the initial medial meshes are generated by the method of Amenta and Bern [1998].

For each model, we show different levels of simplified medial meshes and some reconstruction results. It can be seen that all the spikes are pruned before the number of vertices is reduced to $10,000$, resulting in structurally clean and yet accurate simplified medial meshes. Furthermore, the geometric features of the input models are well preserved at all simplification levels even when there are fewer than 100 medial vertices left. Hence, Q-MAT not only prunes spikes effectively but also reduces the number of mesh vertices drastically while maintaining high approximation accuracy. The simplified medial meshes by Q-MAT meet the aforementioned three requirements for MAT simplification: (1) *structural simplicity*; (2) *geometric accuracy*; and (3) *compactness of representation*. Furthermore, Q-MAT is very fast, due to its use of quadratic error term and adoption of the QEM framework for mesh decimation.

Topology preservation: Q-MAT produces, in most practical settings, a medial mesh of the same homotopy type as that of the input object, but the homotopy type of the object reconstructed from the simplified MAT is not guaranteed to be the same as that of the original object. For example, a change in homotopy happens when extreme simplification causes self-intersection of the reconstructed shape from the simplified medial mesh, as shown in Fig. 20.

Computation time: Table III lists the timing of initialization and timing of using Q-MAT for MAT simplification for various models. Because of its adoption of the QEM framework for efficient maintenance of the quadratic error terms and error minimization, Q-MAT is very fast in comparison with other existing methods.

6. CONCLUSIONS AND FUTURE WORK

We have presented an efficient method, called *Q-MAT*, for computing a structurally simple, compact, and accurate linear approximation of the medial axis transform. The efficacy of Q-MAT is due

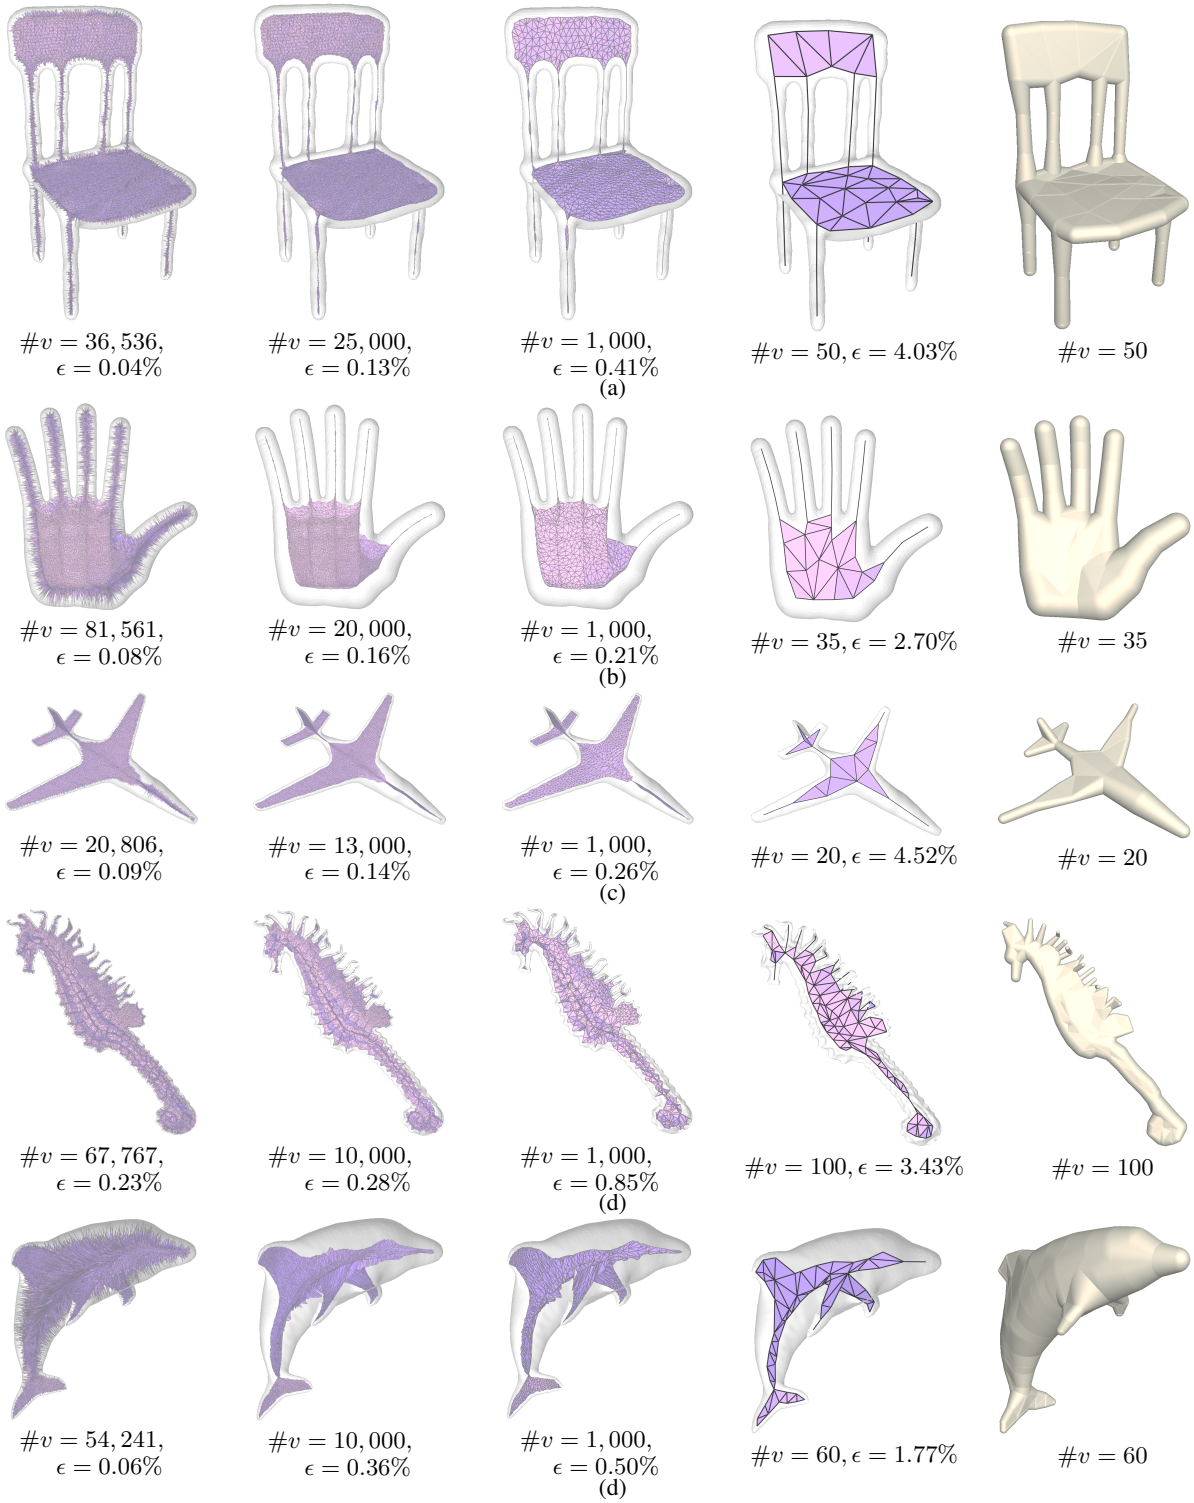

Fig. 18. MAT simplification results by Q-MAT. The first column lists the initial medial meshes. The second column through the fourth column show medial meshes of different levels of simplification. The last column shows the shapes reconstructed from the simplified medial meshes in the fourth column. All the results are computed with $k = 10^{-5}$ in the edge-collapse cost function (Eq. 8).

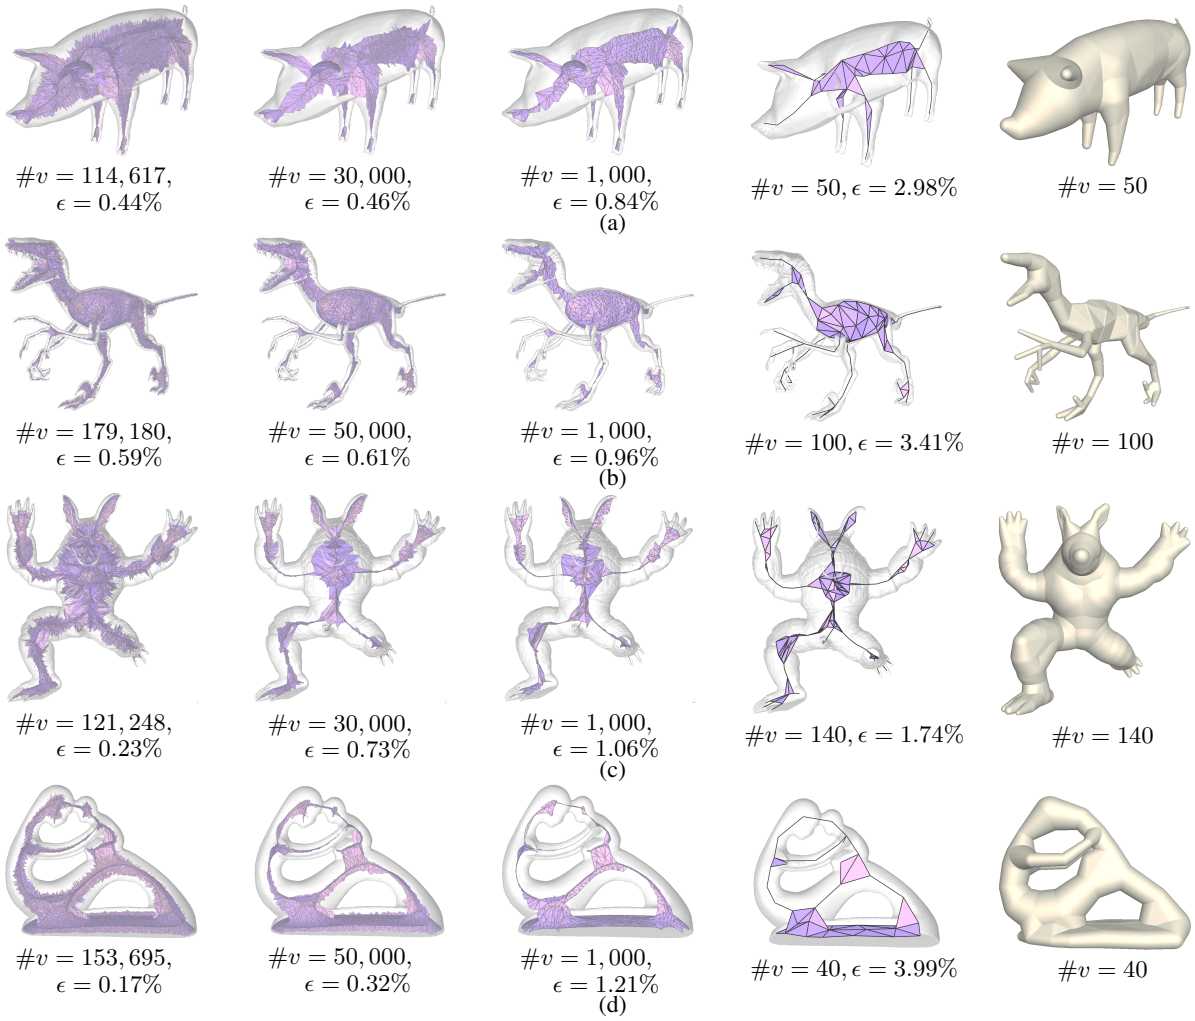

Fig. 19. MAT simplification results by Q-MAT. All the initial medial meshes are generated using the SAT method [Miklos et al. 2010]. The Raptor model is generated using $\delta = 0.006$ and $s = 1.0$, and the others using $\delta = 0.01$ and $s = 1.0$. The first column shows the initial medial meshes. The second column through the fourth column show medial meshes of different levels of simplification. The last column shows the shapes reconstructed from the simplified medial meshes in the fourth column. All the results are computed with $k = 10^{-5}$ in the edge-collapse cost function (Eq. 8).

to two novel metrics we proposed: the *slab quadratic error metric* and the *stability ratio*. The *slab quadratic error metric* measures approximation errors in MAT simplification and the *stability ratio* measures the geometric significance of each edge in the medial mesh in order to distinguish and prune spikes. Q-MAT is also capable of computing an extreme simplification of an input shape to produce a faithful volume approximation represented by a medial mesh with a small number of vertices.

However, the current version of Q-MAT cannot preserve very well sharp features on 3D shapes. Since sharp feature preservation is important in many applications, such as CAD/CAM, further research is needed to address this problem. Another outstanding issue is how to accurately and efficiently convert a medial mesh representation of a 3D object to a boundary surface representation.

ACKNOWLEDGMENTS

We would like to thank Kanglai Qian and Yingya Wei for helping preparing the video and some illustrations, and the anonymous reviewers for valuable suggestions. The models we use in our paper are provided by the database of [Chen et al. 2009] and Aim-at-Shape. This project was partially funded by National Basic Research Program of China (2011CB302400), National Science Foundation of China (61373071, 61272019 and 61332015), and the Research Grant Council of Hong Kong (718311 and 717813). Xiaohu Guo is partially supported by Cancer Prevention & Research Institute of Texas (CPRIT) under Grant No. RP110329, and National Science Foundation (NSF) under Grant Nos. IIS-1149737 and CNS-1012975.

REFERENCES

AMENTA, N. AND BERN, M. 1998. Surface reconstruction by voronoi filtering. In *Proceedings of the Fourteenth Annual Symposium on Com-*

Table III. Computation time of Q-MAT.

Model (#V)	Initialization (ms)	Simplification (ms)
Ant (29,703)	869	4,075
Vase (140,607)	5,754	36,366
Bear (34,703)	1,084	5,603
Venus (74,998)	2,316	12,067
Chair (36,536)	1,051	4,386
Hand (81,561)	2,696	12,216
Plane (20,806)	569	2,530
Dolphin (54,241)	1,689	9,586
Seahorse (67,767)	1,990	10,937
Octopus (39,342)	1,100	5,560
Pig (114,617)	4,744	28,390
Raptor (179,180)	7,542	45,245
Armadillo (121,248)	5,704	32,148
Fertility (153,695)	6,252	40,221
Neptune (95,420)	4,089	23,496

"#V" is the number of vertices in the initial *medial mesh*. "Initialization" includes the time of computing matrices of the initial slab quadric error functions and that of choosing optimal contraction target points. "Simplification" shows the time from the beginning to the end when no edge remains.

putational Geometry. SCG '98. ACM, New York, NY, USA, 39–48.

AMENTA, N., CHOI, S., AND KOLLURI, R. K. 2001. The power crust. In *Proceedings of the sixth ACM symposium on Solid modeling and applications*. ACM, 249–266.

ATTALI, D., BOISSONNAT, J.-D., AND EDELSBRUNNER, H. 2009. Stability and computation of medial axes-a state-of-the-art report. In *Mathematical foundations of scientific visualization, computer graphics, and massive data exploration*. Springer, 109–125.

ATTALI, D. AND MONTANVERT, A. 1996. Modeling noise for a better simplification of skeletons. In *Image Processing, 1996. Proceedings., International Conference on*. Vol. 3. IEEE, 13–16.

BLUM, H. ET AL. 1967. A transformation for extracting new descriptors of shape. *Models for the perception of speech and visual form 19,* 5, 362–380.

CHAZAL, F. AND LIEUTIER, A. 2005. The λ-medial axis. *Graphical Models 67,* 4, 304–331.

CHEN, X., GOLOVINSKIY, A., AND FUNKHOUSER, T. 2009. A benchmark for 3d mesh segmentation. In *ACM Transactions on Graphics (TOG)*. Vol. 28. ACM, 73.

CHOI, S. W. AND SEIDEL, H.-P. 2001. Hyperbolic hausdorff distance for medial axis transform. *Graphical Models 63,* 5, 369–384.

DEY, T. K., EDELSBRUNNER, H., GUHA, S., AND NEKHAYEV, D. V. 1998. Topology preserving edge contraction. *Publ. Inst. Math. (Beograd) (N.S 66*, 23–45.

DEY, T. K. AND ZHAO, W. 2004. Approximate medial axis as a voronoi subcomplex. *Computer-Aided Design 36,* 2, 195–202.

FARAJ, N., THIERY, J.-M., AND BOUBEKEUR, T. 2013. Progressive medial axis filtration. In *SIGGRAPH Asia 2013 Technical Briefs*. ACM, 3.

FOSKEY, M., LIN, M. C., AND MANOCHA, D. 2003. Efficient computation of a simplified medial axis. *Journal of Computing and Information Science in Engineering 3,* 4, 274–284.

GARLAND, M. AND HECKBERT, P. S. 1997. Surface simplification using quadric error metrics. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., 209–216.

JI, Z., LIU, L., AND WANG, Y. 2010. B-mesh: A modeling system for base meshes of 3d articulated shapes. In *Computer Graphics Forum*. Vol. 29. Wiley Online Library, 2169–2177.

MIKLOS, B., GIESEN, J., AND PAULY, M. 2010. Discrete scale axis representations for 3d geometry. *ACM Transactions on Graphics (TOG) 29,* 4, 101.

PIXOLOGIC. 2001. Zbrush.

PIZER, S. M., SIDDIQI, K., SZÉKELY, G., DAMON, J. N., AND ZUCKER, S. W. 2003. Multiscale medial loci and their properties. *International Journal of Computer Vision 55,* 2-3, 155–179.

SIDDIQI, K. AND PIZER, S. M. 2008. Medial representations. *Mathematics, algorithms and applications*.

SUD, A., FOSKEY, M., AND MANOCHA, D. 2005. Homotopy-preserving medial axis simplification. In *Proceedings of the 2005 ACM Symposium on Solid and Physical Modeling*. SPM '05. ACM, New York, NY, USA, 39–50.

SUN, F., CHOI, Y., YU, Y., AND WANG, W. 2013. Medial meshes for volume approximation. *CoRR abs/1308.3917*.

SUN, F., CHOI, Y., YU, Y., AND WANG, W. 2014. Medial meshes: A compact and accurate medial shape representation. *TVCG, submitted*.

THIERY, J.-M., GUY, É., AND BOUBEKEUR, T. 2013. Sphere-meshes: shape approximation using spherical quadric error metrics. *ACM Transactions on Graphics (TOG) 32,* 6, 178.

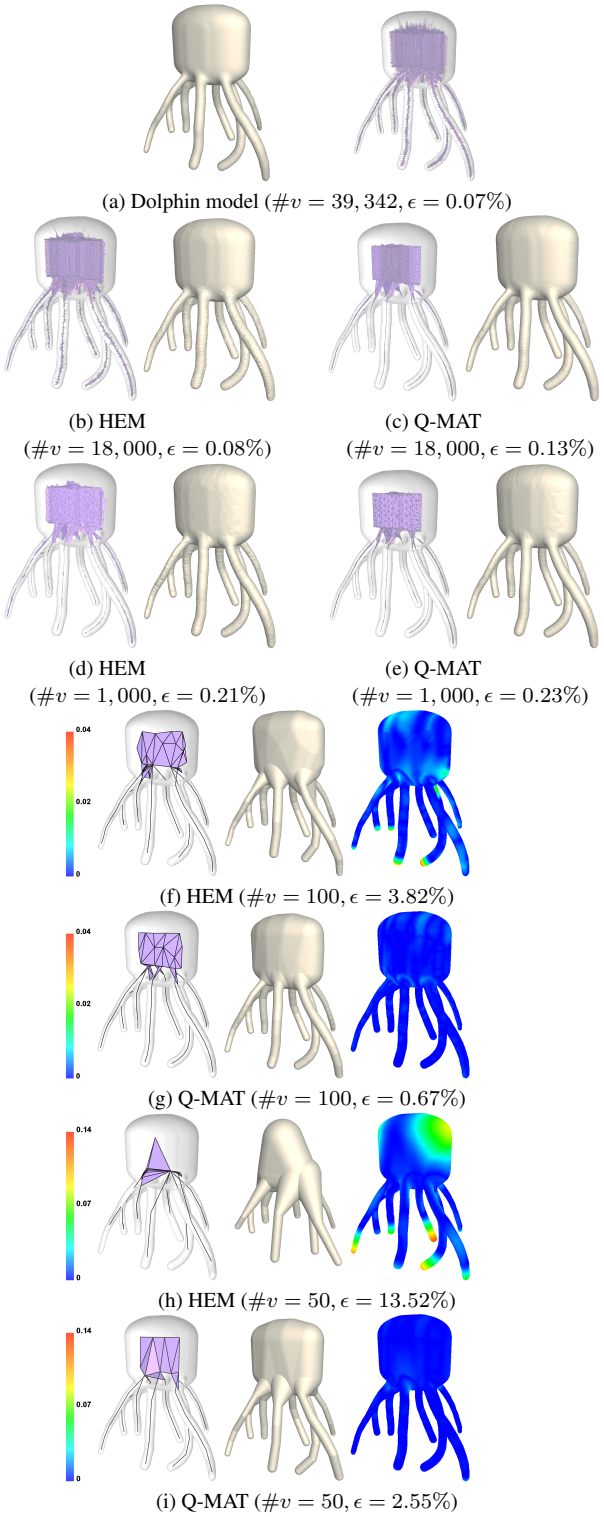

(a) Dolphin model ($\#v = 39,342, \epsilon = 0.07\%$)

(b) HEM ($\#v = 18,000, \epsilon = 0.08\%$)

(c) Q-MAT ($\#v = 18,000, \epsilon = 0.13\%$)

(d) HEM ($\#v = 1,000, \epsilon = 0.21\%$)

(e) Q-MAT ($\#v = 1,000, \epsilon = 0.23\%$)

(f) HEM ($\#v = 100, \epsilon = 3.82\%$)

(g) Q-MAT ($\#v = 100, \epsilon = 0.67\%$)

(h) HEM ($\#v = 50, \epsilon = 13.52\%$)

(i) Q-MAT ($\#v = 50, \epsilon = 2.55\%$)

Fig. 14. Comparison with HEM. (a) The input surface and the initial medial mesh. (b), (d), (f) and (h): Simplification results by HEM. (c), (e), (g) and (i): Simplification results by Q-MAT. Groups (b), (c), (d) and (e), from left to right, the simplified medial mesh and the reconstructed surface from it. Groups (f), (g), (h) and (i) also show the color-coded distributions of errors from the input mesh to the reconstructed surface, respectively.

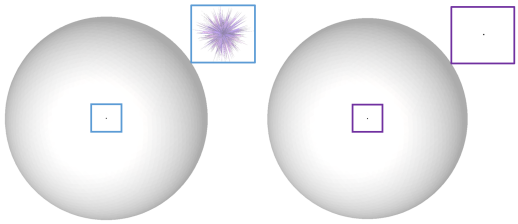

(a) $\#v = 7,157$ (b) $\#v = 1$

Fig. 15. MAT Simplification of Sphere model. (a) The initial medial mesh. (b) The simplified medial mesh.

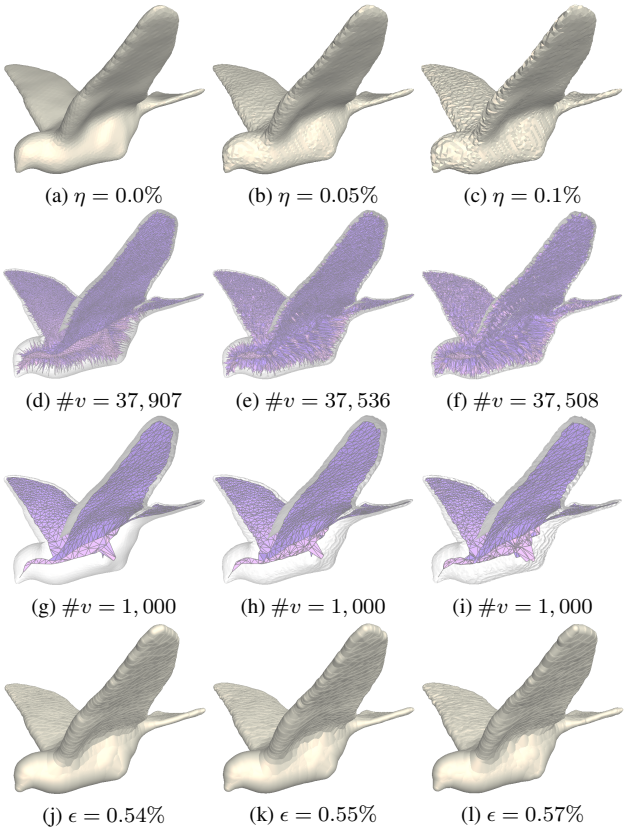

(a) $\eta = 0.0\%$ (b) $\eta = 0.05\%$ (c) $\eta = 0.1\%$

(d) $\#v = 37,907$ (e) $\#v = 37,536$ (f) $\#v = 37,508$

(g) $\#v = 1,000$ (h) $\#v = 1,000$ (i) $\#v = 1,000$

(j) $\epsilon = 0.54\%$ (k) $\epsilon = 0.55\%$ (l) $\epsilon = 0.57\%$

Fig. 16. MAT simplification by Q-MAT with different levels of noises (η) added to the same Bird model, which are generated with random per-vertex displacements expressed in percentages w.r.t. the diagonal length of the model's bounding box. For each noise level we show the initial boundary surface, the initial medial mesh of approximately the same number of vertices, the simplified medial mesh of $1,000$ vertices, and the reconstructed surface from the simplified medial mesh.

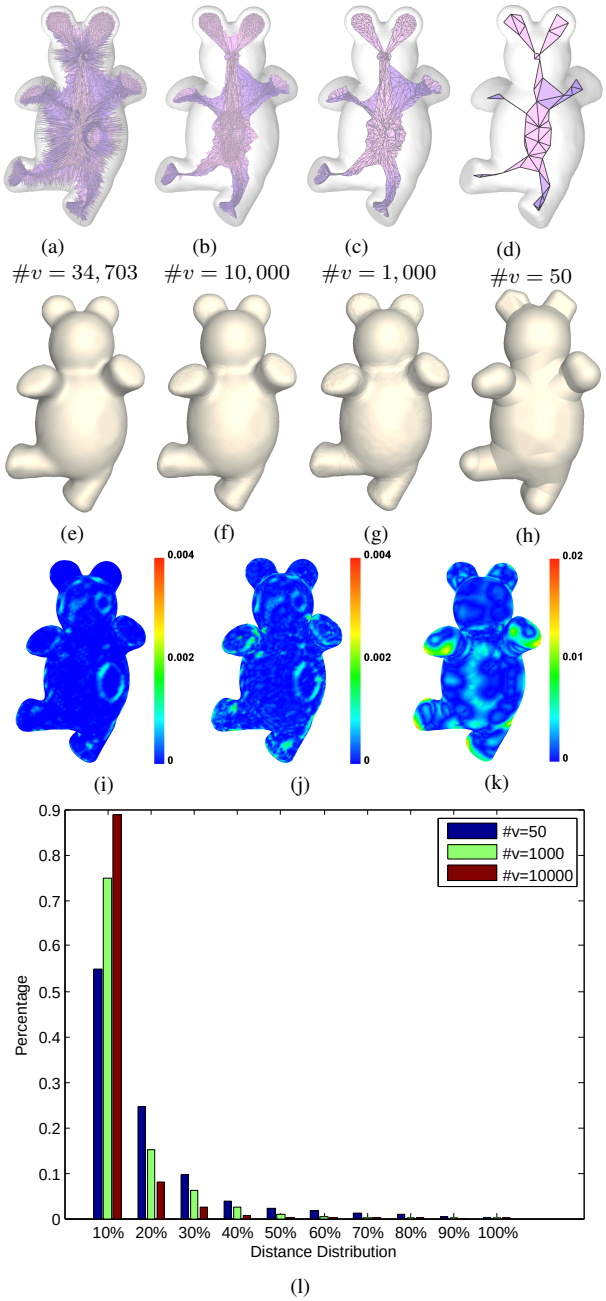

Fig. 17. MAT simplification of the Bear model by Q-MAT. (a) The initial medial axis. (b), (c) and (d): The simplified medial mesh with different number of vertices. (e) The input boundary mesh. (f), (g) and (h): The reconstructed shapes from the medial meshes in (b), (c) and (d), respectively; (i), (j) and (k): the color-coded distribution of the one-sided Hausdorff errors from the input boundary mesh to the reconstructed boundary surfaces in (f), (g) and (h), respectively. (l) shows the distribution of error distance from 10 percentage of the Hausdorff distance to 100 percentage of it.

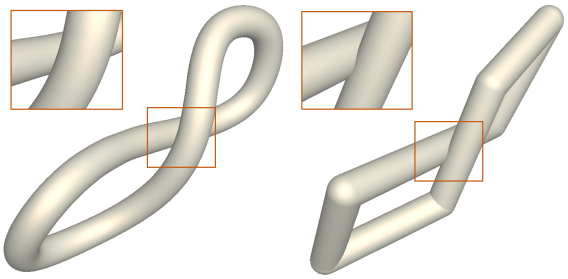

(a) Input mesh (b) $\#v = 6$

Fig. 20. Self-intersection in an extreme simplification situation. (a) The input pipe surface without self-intersection. (b) There is self-intersection in the surface reconstructed from a simplified medial mesh with 6 medial spheres.