Medial Meshes – A Compact and Accurate Representation of Medial Axis Transform

Feng Sun, Yi-King Choi, Yizhou Yu, *Member, IEEE*, and Wenping Wang, *Member, IEEE*

Abstract—The medial axis transform has long been known as an intrinsic shape representation supporting a variety of shape analysis and synthesis tasks. However, for a given shape, it is hard to obtain its faithful, concise and stable medial axis, which hinders the application of the medial axis. In this paper, we introduce the *medial mesh*, a new discrete representation of the medial axis. A medial mesh is a 2D simplicial complex coupled with a radius function that provides a piecewise linear approximation to the medial axis. We further present an effective algorithm for computing a concise and stable medial mesh for a given shape. Our algorithm is quantitatively driven by a shape approximation error metric, and progressively simplifies an initial medial mesh by iteratively contracting edges until the approximation error reaches a predefined threshold. We further demonstrate the superior efficiency and accuracy of our method over existing methods for medial axis simplification.

Index Terms—Medial representation, medial axis simplification, enveloping primitives

1 Introduction

The medial axis of a solid object in $\mathbb{R}^d$, $d = 2$ or 3, is the set of points having at least two closest points on the boundary of the object. In other words, it comprises the centers of the spheres (or circles in 2D) which are contained in the object and touch the boundary of the object at two or more points. These spheres are called the *medial spheres* (or *medial circles* in 2D). The medial axis transform (MAT) consists of a medial axis as the locus of the centers of medial spheres and a radius function encoding the radii of the associated medial spheres. As an intrinsic shape representation, the MAT has proven highly useful for shape analysis and synthesis, such as the approximation, description, recognition and retrieval of shapes as well as topology representation and data reduction of complex models.

Despite its potential utility, the application of the MAT has been impeded by its instability and redundancy. First, the MAT is notoriously sensitive to noisy, small perturbations on the shape boundary, meaning that a slightly noisy shape boundary commonly observed in practice leads to numerous undesired branches of the medial axis. Second, the MAT is in general a continuum comprising infinitely many points without a simple analytical expression. As a result, a medial axis is often approximated with a dense set of sample points, resulting in inefficient representation of the medial axis.

In this work we shall address the following issues: (1) How to simplify a densely sampled MAT of a given 3D object to produce a compact representation with a reduced data size? (2) How to ensure that the resulting approximation represents the given 3D object accurately? In addition, we need to address the instability issue in medial axis computation in the process.

Compared to using the union of many spheres, the MAT can be approximated more accurately and compactly with a 2D non-manifold simplicial complex embedded in 4D, which we call a *medial mesh*, consisting of points, line segments and triangle faces. The vertices of a medial mesh are sampled medial spheres of the MAT, and the linear interpolation of these vertices over line segments or triangles of the medial mesh defines the convex hulls of two or three medial spheres.

We shall present an effective algorithm for computing a stable and compact medial mesh. Our algorithm is based on medial axis simplification with respect to shape approximation error. Our method not only cleans up the topology of the medial axis by pruning unstable branches, but also produces a compact representation by reducing the number of sample points on the medial axis. Analogous to mesh simplification, our algorithm progressively simplifies an initial medial mesh of the input shape by iteratively contracting selected edges until the approximation error reaches a predefined threshold. Using sphere interpolation over the medial axis, this algorithm is capable of drastically reducing the number of vertices and edges in the medial mesh while faithfully representing the original shape. See Fig. 1 for an example. Experiments and comparisons indicate that our simplified medial mesh can achieve the same shape approximation error with orders-of-magnitude fewer primitives than by existing medial axis pruning techniques.

2 Related Work

Medial axis simplification The medial axis transform was first proposed by Blum [1] as a tool for biological applications. It has been proven that the MAT is a complete shape descriptor. The medial representation captures the shape intrinsically by encoding the local thickness and symmetry, and finds applications in shape matching, shape recognition

• *The authors are with the Department of Computer Science, The University of Hong Kong, Hong Kong, China.*
E-mail: {fsun, ykchoi, yzyu, wenping}@cs.hku.hk.

Manuscript received 7 June 2014; revised 15 May 2015; accepted 8 June 2015. Date of publication 28 June 2015; date of current version 3 Feb. 2016.
Recommended for acceptance by S.-M. Hu.
For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.
Digital Object Identifier no. 10.1109/TVCG.2015.2448080

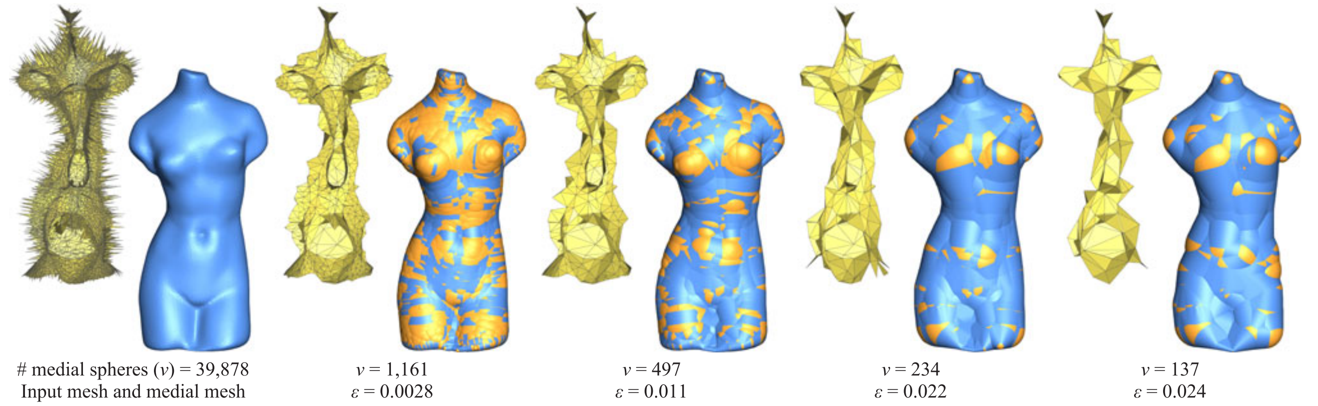

Fig. 1. A series of four progressively simplified medial meshes computed by our method with controlled approximation error. Each represents a shape as the union of enveloping primitives of the corresponding simplified medial mesh comprising spheres (in orange), cones and triangles (in blue). The approximation error (ε), normalized by the diagonal length of the bounding box of the input object, is measured by the one-sided Hausdorff distance from the original shape to the shape represented by the medial axis. The time taken to generate the four medial meshes in this series are 200, 226, 250 and 264 seconds, respectively.

and shape retrieval [2], [3], [4]. However, the medial axis is inherently unstable, that is, a small perturbation to the shape boundary may introduce a large change of its medial axis. In graphics, shapes are often represented by their boundaries as triangle meshes. Typically, the exact medial axis of a shape represented by a boundary mesh has many undesired branches [5], [6], and therefore is not suitable for further applications. To overcome the instability of the medial axis, a number of methods have been proposed for removing the unstable branches, a process often referred to as medial axis pruning. We review below several prevailing approaches for medial axis simplification.

Angle-based filtering. Recall that a point on the medial axis has at least two closest points on the object boundary. Angle-based methods [7], [8], [9], [10] compute the angle at the point subtended by its closest points. (For a medial point with multiple closest points on the boundary, the largest angle spanned by any two of these closest point is taken.) Those vertices on the medial axis with their subtended angles being smaller than a user-specified threshold are removed and those remaining points are kept to define the filtered medial axis. Although the angle criterion preserves local features, angle-based filtering often yields a simplified medial axis with a different topology from the input one [11].

The λ medial axis. The λ medial axis is another method for removing unstable branches. The criterion of the λ medial axis [12], [13] is the circumradius of the closest boundary points of a medial point. Any medial point is removed if its circumradius thus defined is smaller than a given threshold λ. However, this circumradius criterion does not work well on shapes with features at different scales. Small values of λ cannot remove noise near large-scale features while increasing λ would eliminate small-scale features, potentially leading to a different homotopy type, though it is proven that λ medial axis preserves the homotopy for small λ [12]. As a result, there may exist large discrepancies between the original shape boundary and the boundary of the shape represented by the simplified medial axis.

Scale axis transformation. Miklos et al. [11] propose a method based on the scale axis transformation (SAT). For a given medial axis, all medial spheres are first scaled by a factor s larger than 1. A scaled medial sphere is removed if it is contained in another scaled medial sphere. Then, the union of all remaining medial spheres generates a new shape, whose medial axis is further subject to homotopy-preserving angle filtering. The final result is obtained by shrinking all medial spheres of the simplified medial axis by the factor $1/s$. This SAT method often generates results better than previous techniques. However, it does not preserve shape homotopy as it may fill in narrow gaps or small holes by the first dilating step and therefore change the homotopy of the medial axis. A variant with accelerated computation and improved homotopy preservation is proposed in [14].

Feature-based simplification. Another approach to medial axis simplification is taken by Tam and Heidrich [15] to achieve high-level feature-based shape simplification. The medial axis is decomposed into manifold sheets as parts, each of which corresponds to a feature of the input shape. The parts are then pruned based on significance measures using triangle count and the volume of each part. While the method removes insignificant shape features and manifold sheets with volumes smaller than a specified volume threshold, it does not reduce the number of medial points on the remaining manifold sheets; that is, the geometric complexity of these parts remains high. In contrast, our simplification method yields a medial axis that is both stable and compact, since we reduce the number of sample points on the medial axis as much as possible while maintaining the required approximation accuracy.

All the above methods do not use the geometric approximation error to control the simplification process. As a consequence, a simplified medial axis produced by these methods often fails to represent the original shape faithfully. In contrast, our method maintains explicitly the geometric approximation error of each vertex of the simplified mesh and therefore ensures the approximation quality as specified by the user.

Sphere interpolation Note that all previous pruning methods approximate an input shape using a discrete set of individual spheres. Such a discretization simplifies algorithm design but gives rise to artifacts or large shape approximation errors. Previous works propose to approximate the medial axis as a Voronoi sub-complex [9], [16] by including both Voronoi edges and Voronoi faces.

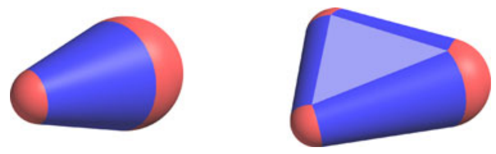

Fig. 2. Left: The enveloping primitive of a medial edge. Right: The enveloping primitive of a medial face.

Meanwhile, linear interpolation of spheres has been used [17], [18], [19], [20], [21] to obtain a swept sphere volume for different applications. The medial mesh we proposed uses linear interpolations of spheres to approximate an input shape to achieve both visual and numerical fidelity.

M-rep A closely related work of MAT is M-rep by Pizer et al. [22], which proposes the compact spline approximation to the medial axis of 3D objects for shape analysis in medical imaging. However, M-rep considers only the special case that the 3D object is simple enough to allow its medial axis to be approximated by a single patch of tensor product B-spline surface. Furthermore, it assumes that this tensor product B-spline surface patches is manually specified and does not consider how to extract such a spline representation automatically for a given object. The idea of the medial mesh is inspired by the M-rep, but our goal is to compute an accurate and compact medial axis approximation of *any* 3D object based on piecewise linear interpolation of medial spheres.

Straight skeletons The straight skeleton [23] is also related to the medial axis, since both provide an $(n-1)$-dimensional representation of an n-dimensional shape. However, straight skeletons are defined only for polygons, rather than general domains with curved boundaries. The straight skeleton of a polygon is the trace of vertices when all edges are moved inwards parallel to themselves with the same speed; it comprises only line segments, thus finding applications in domain decomposition [24], paper origami [25] and shape morphing [26]. In contrast with medial axes, straight skeletons have branches corresponding to both convex bumps and concave dents of the polygons. While branches due to convex bumps provide meaningful and intuitive representation of object parts, branches induced by dents are often hard to work with. The medial axis is therefore considered more suitable for shape representation in most applications. The medial mesh proposed in this paper, though also a linear structure as the straight skeleton, is an approximation to the medial axis for compact and accurate shape representation.

Manifold/Skeleton approximation of medial axes To accommodate for applications in shape approximation and deformation, Yoshizawa et al. [27] approximate the medial axis of an object by a manifold mesh. They iteratively contract the input boundary mesh towards the medial axis by applying a bi-Laplacian flow to the mesh. A manifold mesh is finally obtained to approximate the medial axis. Note that a medial sheet of the medial axis is approximated by two close sheets of triangles on each side of the mesh. This method is further generalized by Au et al. [28] to generate a curved skeleton of the input shape as an approximate medial axis. Liu et al. [29] generate a stable skeleton of an object through thinning the noisy skeleton by measuring its distance to the boundary in a discrete manner.

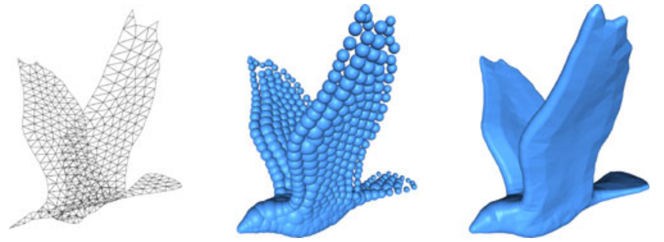

Fig. 3. From left to right: A medial mesh of a bird shape with 400 medial vertices. The shape represented by the union of the 400 sample medial spheres. The shape represented by the union of the enveloping primitives of the medial mesh.

Recently Thiery et al. [30] propose a volume simplification method. They define a zero-radius sphere on vertices of the input mesh and then iteratively contract the mesh edges and modify the sphere radius on the vertices to approximate the input mesh by the envelope of the spheres. The sphere mesh thus obtained is often far away from the medial axis. Although the envelope of the sphere mesh may approximate the input mesh to some extent, it does not fall in the category of medial axis simplification which is the main focus of this paper. Similar to the simplification methods, the above methods do not enforce an explicit approximation error control, and their results are only good for shape abstraction but not for accurate shape representation in general.

3 Medial Mesh Representation

The MAT of a shape in $\mathbb{R}^3$ is in general a two-dimensional CW-complex embedded in $\mathbb{R}^4$, since each point of the medial axis consists of the coordinates of its 3D position plus the radius of its associated medial sphere. We approximate the medial axis surface with a 2D simplicial complex, which we call the *medial mesh*. A vertex of a medial mesh is called a *medial vertex* and is a point $\mathbf{v} = (\mathbf{p}, r) \in \mathbb{R}^4$, where $\mathbf{p} \in \mathbb{R}^3$ is the position of the vertex and r its associated radius value. An edge $e = \{\mathbf{v}_1, \mathbf{v}_2\}$ of the medial mesh is called a *medial edge*, represented by $(1-t)\mathbf{v}_1 + t\mathbf{v}_2, t \in [0,1]$, a convex interpolation of its end points $\mathbf{v}_1$ and $\mathbf{v}_2$. Likewise, a *medial face* is given by $f = \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$, the convex combination of three medial points, i.e., $a_1\mathbf{v}_1 + a_2\mathbf{v}_2 + a_3\mathbf{v}_3$, with $a_1, a_2, a_3 \geq 0$ and $a_1 + a_2 + a_3 = 1$.

Each discrete element of the medial mesh corresponds to an *enveloping primitive* in $\mathbb{R}^3$: a medial vertex being a medial sphere with center $\mathbf{p}$ and radius r, and a medial edge or face being the convex hull of the medial spheres at their respective vertices (Fig. 2). A 3D object is then represented by the *enveloping volume* of the medial mesh, i.e., the union of all the enveloping primitives (Fig. 3). The boundary surface of the enveloping volume is then called the *enveloping surface*. The analogue of the medial mesh for 2D shapes is a graph $(\mathcal{V}, \mathcal{E})$ consisting of a set $\mathcal{V}$ of medial vertices and a set $\mathcal{E}$ of medial edges connecting medial vertices. Here the interpolation of medial spheres for a medial mesh in 3D is replaced by the interpolation of medial circles in 2D.

4 Medial Mesh Computation

We now present an effective algorithm for computing a compact medial mesh that accurately approximates a given

3D shape, satisfying a user-specified error threshold. We shall first discuss how the approximation error is defined and evaluated, and then introduce the steps of our algorithm in detail and discuss how the issues of medial axis stability and homotopy preservation are addressed.

4.1 Approximation Error

In order to gauge approximation error during medial axis simplification, we need to evaluate the approximation error at each step of simplification. While it is natural of think of using the Hausdorff distance between the original shape and the shape represented by a medial mesh as a faithful measure of the approximation error, evaluating the Hausdorff distance is prohibitively costly, which would make the method too inefficient to be practical. Hence, we instead evaluate the one-sided Hausdorff distance from the boundary surface of the input shape to the boundary surface of the shape represented by the simplified medial mesh. We argue that this suffices because: (1) We start with a medial mesh that well approximates an object and our simplification preserves homotopy of the medial mesh at each step (Section 4.2.2); (2) The enveloping surface of the medial mesh is in general simpler than the original surface and the one-sided Hausdorff distance provides a good approximation to the true Hausdorff distance. Note that, while the one-sided Hausdorff distance is used during simplification, for validation purposes the true Hausdorff errors are computed and reported as the approximation error for all the final simplified results presented in this paper. A detailed discussion on MAT simplification and the Hausdorff distance measure can be found in [31].

Let $\mathcal{M} = \{m_i\}$ be a medial mesh approximating an object $\mathcal{S} \subset \mathbb{R}^3$, where m_i denotes a medial vertex, edge or face, in $\mathcal{M}$. Let $E(m_i)$ be the enveloping primitive given by m_i and $E(\mathcal{M}) = \cup E(m_i)$ be the enveloping volume represented by $\mathcal{M}$. The approximation error of $\mathcal{M}$ to $\mathcal{S}$ is then defined as:

$$\xi(\mathcal{M}, \mathcal{S}) = \sup\{d\big(\mathbf{q}, \partial E(\mathcal{M})\big) \,|\, \mathbf{q} \in \partial\mathcal{S}\}, \tag{1}$$

with $d(\mathbf{q}, \partial E(\mathcal{M}))$ being the distance of a point $\mathbf{q} \in \partial\mathcal{S}$ to the enveloping surface $\partial E(\mathcal{M})$ given by

$$d(\mathbf{q}, \partial E(\mathcal{M})) = \min_i\{d\big(\mathbf{q}, \partial E(m_i)\big) \,|\, m_i \in \mathcal{M}\},$$

where $d\big(\mathbf{q}, \partial E(m_i)\big)$ is the distance of a point $\mathbf{q}$ to the surface of an enveloping primitive $E(m_i)$. The computation of $d\big(\mathbf{q}, \partial E(m_i)\big)$ in given in the Appendix.

We also define the set of *governing primitives* of a point $\mathbf{q} \in \partial S$ as

$$\mathrm{Gov}_{\mathcal{M}}(\mathbf{q}) = \{m_i \in \mathcal{M} \,|\, d\big(\mathbf{q}, \partial E(m_i)\big) = d\big(\mathbf{q}, \partial E(\mathcal{M})\big)\},$$

that is, $\mathrm{Gov}_{\mathcal{M}}(\mathbf{q})$ contains the simplices of the medial mesh that are closest to a point $\mathbf{q}$ on the boundary of the input object. The set of *affiliates* of a simplex m_i of $\mathcal{M}$ is defined to comprise those surface points governed by m_i:

$$\mathrm{Aff}_{\mathcal{M}}(m_i) = \{\mathbf{q} \,|\, \mathbf{q} \in \partial S, m_i \in \mathrm{Gov}_{\mathcal{M}}(\mathbf{q})\}.$$

4.2 The Algorithm

Our method performs medial axis simplification with explicit maintenance and control of the shape approximation error. It has two main stages: (1) *initialization*, which aims to obtain a medial mesh of an input shape to start with; and (2) *simplification*, which progressively simplifies the initial medial mesh by iterative edge contraction until the approximation error reaches a predefined threshold.

4.2.1 Initialization

Our algorithm first computes an initial medial mesh for an input shape $\mathcal{S} \subset \mathbb{R}^3$ using the Voronoi-based approach [32]. We first obtain an ε-sample $\mathcal{P}$ on the boundary of $\mathcal{S}$, in which each point $x \in \mathcal{P}$ is of a distance less than $\varepsilon f(x)$ from another point in $\mathcal{P}$, where $f(x)$ is the distance of $x \in \partial\mathcal{S}$ to its medial axis (i.e., the local feature size of $\mathcal{S}$ [32]). We adopt $\varepsilon \leq 0.4$ as suggested by Dey and Zhao [9]. We then compute the Voronoi diagram of $\mathcal{P}$, which is a discrete counterpart of the medial axis, since any point on a Voronoi cell also has at least two closest points among the boundary samples. The set of Voronoi vertices lying inside $\mathcal{S}$, denoted $\mathcal{V}$, are then chosen to approximate the medial axis. The connectivity among the vertices in $\mathcal{V}$ as a subset of the Voronoi vertices of $\mathcal{P}$ is inherited to form a 2D non-manifold mesh. This is taken as the initial medial mesh, which is typically a noisy and dense representation of the medial axis of the input 3D shape.

Now, since $\partial\mathcal{S}$ is well sampled by $\mathcal{P}$, the approximation error of a medial mesh $\mathcal{M}$ to $\mathcal{S}$ in Eq. (1) becomes

$$\xi(\mathcal{M}, \mathcal{S}) = \max\{d\big(\mathbf{p}, \partial E(\mathcal{M})\big) \,|\, \mathbf{p} \in \mathcal{P}\}.$$

Although for a 2D shape the Voronoi vertices of its boundary sample points converges to the medial axis of the shape as the sampling density of $\mathcal{P}$ approaches infinity [33], the same does not extend to the 3D case, due to the existence of slivers close to the shape boundary in the Delaunay triangulation, which is dual to the Voronoi diagram. It has been shown that a subset of Voronoi vertices, called poles, converge to the medial axis [8]. Although our method starts with all interior Voronoi vertices, we will show in Section 4.3 that it can effectively remove those Voronoi vertices corresponding to the circumcenters of the slivers which are far from the medial axis, as well as noisy "spikes" of the medial axis that contribute little to the definition of the input shape.

4.2.2 Simplification

Next we shall apply an edge-contraction scheme to simplify the initial medial mesh. Our aim is to remove redundant geometries to achieve compactness, and to remove spikes to achieve stability of the medial mesh. The basic idea is simple—like other previous mesh simplification schemes based on edge contraction, we iteratively contract selected edges of the medial mesh $\mathcal{M}$ to simplify it progressively until the approximation error of the medial mesh to the given shape $\mathcal{S}$ reaches a user-specified threshold. However, because of the unique properties of the medial mesh, we need to address some new issues that are not encountered in conventional mesh simplification methods developed for 2D mesh surfaces in 3D space.

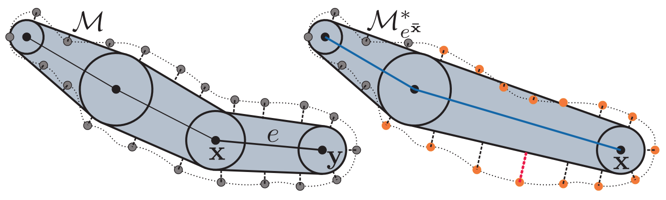

Fig. 4. A 2D illustration for the post-contraction error. The post-contraction error, $\epsilon_{\mathcal{M}}(e^{\bar{\mathbf{x}}})$, of merging vertex $\mathbf{x}$ to vertex $\mathbf{y}$ of the edge e in the medial mesh $\mathcal{M}$ in (a) that results in a new medial mesh $\mathcal{M}^*_{e^{\bar{\mathbf{x}}}}$ in (b). The orange sample points on the shape boundary are affiliated to the 1-ring neighboring edges of $\mathbf{x}$ in $\mathcal{M}$ and the maximum shape approximation error among these points is taken as $\epsilon_{\mathcal{M}}(e^{\bar{\mathbf{x}}})$ (in red). If the merge is realized, the post-contraction errors of the blue medial edges need to be re-evaluated, as they are within the 2-ring neighborhood of $\mathbf{x}$.

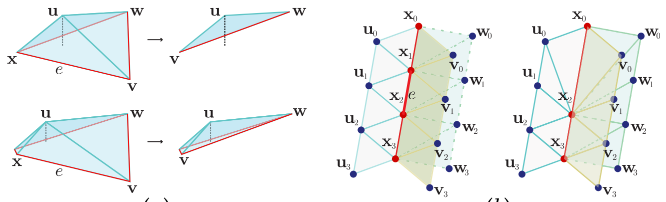

Fig. 5. Manifold preservation. (a) The red polygon on the left is the open boundary of a disk region formed by the medial faces. Contracting the edge e by merging $\mathbf{v}$ to $\mathbf{x}$ results in a non-manifold edge (top) and a *pseudo* nonmanifold structure (bottom). Note the sharp change in the normal direction for the face $\mathbf{uvw}$. (b) Non-manifold vertex $\mathbf{x}_1$ can only merge to the neighboring non-manifold vertices $\mathbf{x}_0$ and $\mathbf{x}_2$ but not the other manifold vertices to retain seams. Merging from a manifold vertex ($\mathbf{v}_1$) to a non-manifold vertex ($\mathbf{x}_1$) is however allowed. Left: Before contracting edge $e = \mathbf{x}_1\mathbf{x}_2$, by merging $\mathbf{x}_1$ to $\mathbf{x}_2$. Right: After edge contraction.

Edge post-contraction error We first define the edge post-contraction errors which govern the order in which the edge are contracted. An edge contraction generally leads to a larger approximation error and there is only a local region on $\partial\mathcal{S}$ near the contracted edge that contributes to this increased error. Given a medial mesh $\mathcal{M}$, we maintain for every possible edge contraction a *post-contraction error*, which measures the local shape approximation error caused by the contraction of an edge. An edge is contracted by merging either one of its vertices to another, and there are therefore two possible ways of merging, each entailing its own contraction error. Let $e^{\bar{\mathbf{x}}}$ denote an edge contraction by merging vertex $\mathbf{x}$ to vertex $\mathbf{y}$ for an edge $e = (\mathbf{x}, \mathbf{y})$, which results in a new medial mesh $\mathcal{M}^*_{e^{\bar{\mathbf{x}}}}$. Then the post-contraction error is defined as:

$$\begin{aligned}\epsilon_{\mathcal{M}}(e^{\bar{\mathbf{x}}}) = \max\{d\big(\mathbf{p}, \partial E(\mathcal{M}^*_{e^{\bar{\mathbf{x}}}})\big) \,|\, \mathbf{p} \in \mathrm{Aff}_{\mathcal{M}}(m_i), \\ m_i \text{ is an edge or a face in } \mathcal{M} \text{ incident to } \mathbf{x}\}.\end{aligned} \tag{2}$$

In other words, those points in $\mathcal{P}$ which are the affiliates of the medial faces incident to $\mathbf{x}$ contribute to the contraction error. Fig. 4 explains how the post-contraction error of an edge is evaluated in a 2D setting. The contraction from $\mathbf{y}$ to $\mathbf{x}$ induces a post-contraction error $\epsilon_{\mathcal{M}}(e^{\bar{\mathbf{y}}})$ similarly.

Edge contraction Algorithm 1 details the simplification steps. The post-contraction error for every possible edge contraction in the initial medial mesh $\mathcal{M}_0$ is first computed using Eq. (2) and is maintained by a priority queue $\mathcal{Q}_0$. The function *ValidEdge*$(\mathcal{M}_k, \mathcal{Q}_k, \varepsilon)$ on line 3 then iteratively retrieves an edge contraction $e^{\bar{\mathbf{x}}}$ with the smallest contraction error from $\mathcal{Q}_k$, which is only returned if it also: (1) satisfies that the contraction error $\epsilon_{\mathcal{M}}(e^{\bar{\mathbf{x}}})$ is no greater than the error threshold ε; *and* (2) satisfies the two topological conditions for (i) preserving the medial mesh manifold, and (ii) avoiding self-intersection in the enveloping surface. The topological conditions will be discussed later in this section.

Once a valid edge contraction $e^{\bar{\mathbf{x}}}$ is identified, *EdgeContract*$(\mathcal{M}_k, e^{\bar{\mathbf{x}}})$ on line 4 realizes the contraction by merging the vertices of the edge e, and collapsing the medial faces incident to e. This results in a new medial mesh $\mathcal{M}_{k+1}$. Since the change is local to the 1-ring neighborhood of vertex $\mathbf{x}$ in $\mathcal{M}$, the governing primitives of all surface sample points in $\mathcal{P}$ that are the affiliates of this 1-ring neighborhood need to be updated (lines 5-7). Now, the post-contraction error in the 2-ring neighborhood is affected (as the error is defined over the 1-ring neighborhood of an edge) and hence is updated (Fig. 4). (The functions $\text{1-ring}_{\mathcal{M}}(\mathbf{x})$ and $\text{2-ring}_{\mathcal{M}}(\mathbf{x})$ stand for the 1-ring and 2-ring neighborhood of a vertex $\mathbf{x}$ in $\mathcal{M}$, respectively.) When there is no more valid edge to contract (i.e., *ValidEdge* in line 3 returns nil), the algorithm terminates and outputs the final simplified medial mesh $\mathcal{M}_f$. Due to the error control by the function *ValidEdge*, the approximation error of $\mathcal{M}_f$ to $\partial\mathcal{S}$ sampled by $\mathcal{P}$ is guaranteed to be no greater than the user threshold ε.

Algorithm 1. Medial Mesh Simplification

Input: Surface samples $\mathcal{P}$ of a shape $\mathcal{S}$, initial medial mesh $\mathcal{M}_0$, error threshold ε
Output: Final medial mesh $\hat{\mathcal{M}}$ with $\xi(\hat{\mathcal{M}}, \mathcal{S}) \leq \varepsilon$

1: $k \leftarrow 0$
2: Initialize priority queue $\mathcal{Q}_k$ to store all possible edge contractions $e^{\bar{\mathbf{x}}}$ ordered by post-contraction errors $\epsilon_{\mathcal{M}_k}(e^{\bar{\mathbf{x}}})$
3: **while** $\big((e^{\bar{\mathbf{x}}}, \mathcal{Q}_{k+1}) \leftarrow$ ***ValidEdge***$(\mathcal{M}_k, \mathcal{Q}_k, \varepsilon)\big) \neq$ *nil* **do**
4: $\quad\mathcal{M}_{k+1} \leftarrow$ ***EdgeContract***$(\mathcal{M}_k, e^{\bar{\mathbf{x}}})$
5: $\quad$**for all** $\mathbf{p} \in \mathrm{Aff}_{\mathcal{M}}\big(\text{1-ring}_{\mathcal{M}}(\mathbf{x})\big)$ **do**
6: $\quad\quad$Update $\mathrm{Gov}_{\mathcal{M}_{k+1}}(\mathbf{p})$, $\mathrm{Aff}_{\mathcal{M}_{k+1}}(\mathrm{Gov}_{\mathcal{M}_{k+1}}(\mathbf{p}))$ and $\mathrm{Aff}_{\mathcal{M}_{k+1}}(\mathrm{Gov}_{\mathcal{M}_k}(\mathbf{p}))$
7: $\quad$**end for**
8: $\quad$**for all** edges $e' = (\mathbf{s}, \mathbf{t}) \in \text{2-ring}_{\mathcal{M}_{k+1}}(\mathbf{x})$ **do**
9: $\quad\quad$Re-evaluate $\varepsilon_{\mathcal{M}_{k+1}}(e'^{\bar{\mathbf{s}}})$ and $\varepsilon_{\mathcal{M}_{k+1}}(e'^{\bar{\mathbf{t}}})$
10: $\quad\quad$Update $\mathcal{Q}_{k+1}$
11: $\quad$**end for**
12: $\quad k \leftarrow k + 1$
13: **end while**
14: $\hat{\mathcal{M}} \leftarrow \mathcal{M}_k$

Preserving medial mesh homotopy When removing unstable spikes and simplifying the geometries of the medial mesh, it is important that the topology of the medial mesh at each edge contraction step be preserved (e.g., a disk region is simplified into a disk region). We employ the *link conditions* [34] for ensuring topology-preserving edge contraction, so that the function *ValidEdge* will help invalidate an edge contraction that violates the link conditions. Furthermore, an edge contraction may lead to an abrupt geometric change (although the topology remains the same) and thus a *pseudo* nonmanifold structure (Fig. 5a). We check the normals of the involving medial faces before and after an edge contraction, and invalidate an edge contraction if there is any sharp change in the normal directions.

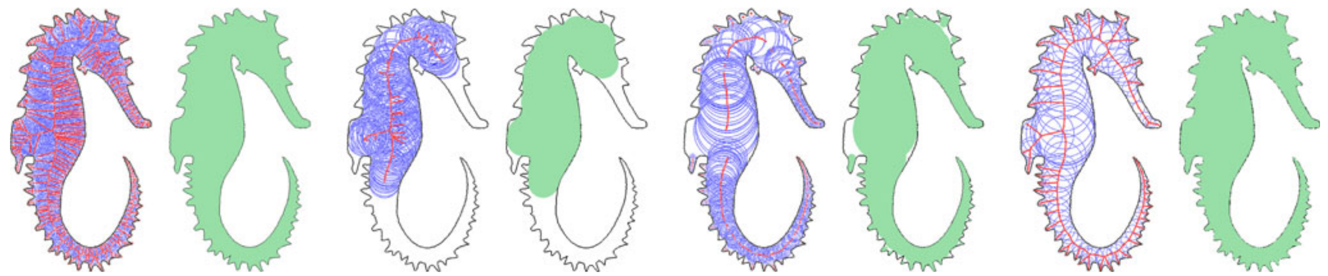

(a) **Original** $v = 5,531,\ \varepsilon = 0.0$ (b) λ **medial axis** $v = 294,\ \varepsilon = 3.0e{-1}$ (c) **Angle-based** $v = 1,429,\ \varepsilon = 3.9e{-2}$ (d) **Medial mesh** $v = 194,\ \varepsilon = 1.6e{-3}$

Fig. 6. A comparison of medial axis simplification methods. The original medial axis or the filtered medial axis (red), the medial circles (blue) and the shapes (green) represented by medial axes of different methods are shown. The medial mesh achieves an accurate approximation with very few medial primitives. (v: number of medial points, ε: approximation error)

Now we consider another type of important medial features, called *seams*. Seams are where large medial sheets meet and are hence made up of non-manifold vertices of the medial mesh. To preserve seams and protect important medial features, we impose a necessary topological condition to validate an edge contraction. Specifically, we do *not* allow a non-manifold vertex to merge to a neighboring manifold vertex. Furthermore, we do *not* allow an endpoint of a seam to merge to an interior point of the seam, according to [34]. However, a manifold vertex is allowed to merge to an adjacent non-manifold vertex. To enable edge simplification along seams, adjacent non-manifold vertices are permitted to merge (Fig. 5b). Note that a non-manifold edge may become a manifold edge if some of its adjacent manifold sheets are pruned.

Handling enveloping surface self-intersection Without proper measures, the medial mesh simplification algorithm may create enveloping primitives that encompass some volume outside of the original shape $\mathcal{S}$, which may cause self-intersection of the shape represented by the medial mesh. We therefore would like to detect and avoid such edge contractions. To this end, we maintain for each surface sample point $\mathbf{p} \in \mathcal{P}$ a *local outgrow bound*, *lob*($\mathbf{p}$), which is defined as the distance of $\mathbf{p}$ to the external medial axis of $\mathcal{P}$. It is analogous to the local feature size [32] which is defined with respect to the internal medial axis instead. Let $d^*(\mathbf{p}, \partial E(M))$ be the signed distance from $\mathbf{p}$ to the enveloping surface of a medial mesh $\mathcal{M}$, with a negative value for points inside $\partial E(M)$. The *lob* can then be used to dictate the maximal outward extension at a point: if a potential edge contraction results in $d^*(\mathbf{p}, \partial E(M)) < lob(\mathbf{p})$ for some $\mathbf{p} \in \mathcal{P}$, the edge contraction would be considered invalid. This checking is local for those sample points affiliated to the affected enveloping primitives in the 1-ring neighborhood of the edge. While we do not observe any local self-intersection in our experiments, it must be noted that the above treatment does not guarantee a self-intersection free enveloping surface, due to the discretization of *lob* and the one-sided Hausdorff distance evaluation of d^*.

These two measures on preserving medial mesh homotopy and avoiding enveloping surface self-intersection serve to preserve the input shape homotopy in practice.

4.3 Stability

The intrinsic instability of MAT refers to the phenomenon that a small perturbation of the boundary of an object may result in a long spike of the medial axis. For a medial mesh, the removal of a medial vertex corresponding to such small perturbation on the boundary does not lead to much change in the shape boundary, therefore yielding only quite small approximation error. As a result, these medial vertices will be merged to their adjacent medial vertices. The same is true of those Voronoi vertices that are the circumcenters of the slivers in the dual Delaunay triangulation. These Voronoi vertices are generally close to the boundary surface, and that the radii of the corresponding medial spheres are small. Again, the removal of these medial spheres would not incur much approximation error. Indeed, we observe experimentally that these medial spheres and the noisy spikes on the medial mesh are removed at early iterations. See Fig. 1 for an example.

5 Experimental Results

In this section, we present the results of our method and its comparison with several existing methods on 3D models of a wide variety of shapes, complexities and topologies. All experiments are performed on a Windows 7 workstation with an Intel i7 CPU and 12 GB main memory. In order to demonstrate the superior simplification ability of our method, we opt to use the Voronoi diagram of sample points on the input boundary surface as the initial medial axis which is typically highly unstable, although our algorithm can take as input the filtered medial axis computed by methods such as the power crust [8]. We compute the Voronoi diagram by using the CGAL package *3D Triangulations* [35] to compute the Delaunay triangulation and then taking its dual. The conversion from a medial mesh to a triangle mesh as a boundary representation of the shape represented by the medial mesh is carried out using the CGAL package *3D Skin Surface Meshing* [36]. Our algorithm is fast and has small memory footprint. For a typical 3D model with 4K vertices, our algorithm generates a simplified medial mesh of 2K medial vertices with approximation error 0.001 (relative to the normalized diagonal of the bounding box of the input shape) using around 20 seconds with 130 MB memory usage.

Comparison with criterion based methods We first compare our method to two existing medial axis pruning methods, the λ medial axis [12] and the angle-based [10] methods, using a 2D seahorse shape shown in Fig. 6. The λ medial axis method performs filtering using the circumradius criterion and completely removes the head and tail of the seahorse while the noise on the trunk still remains,

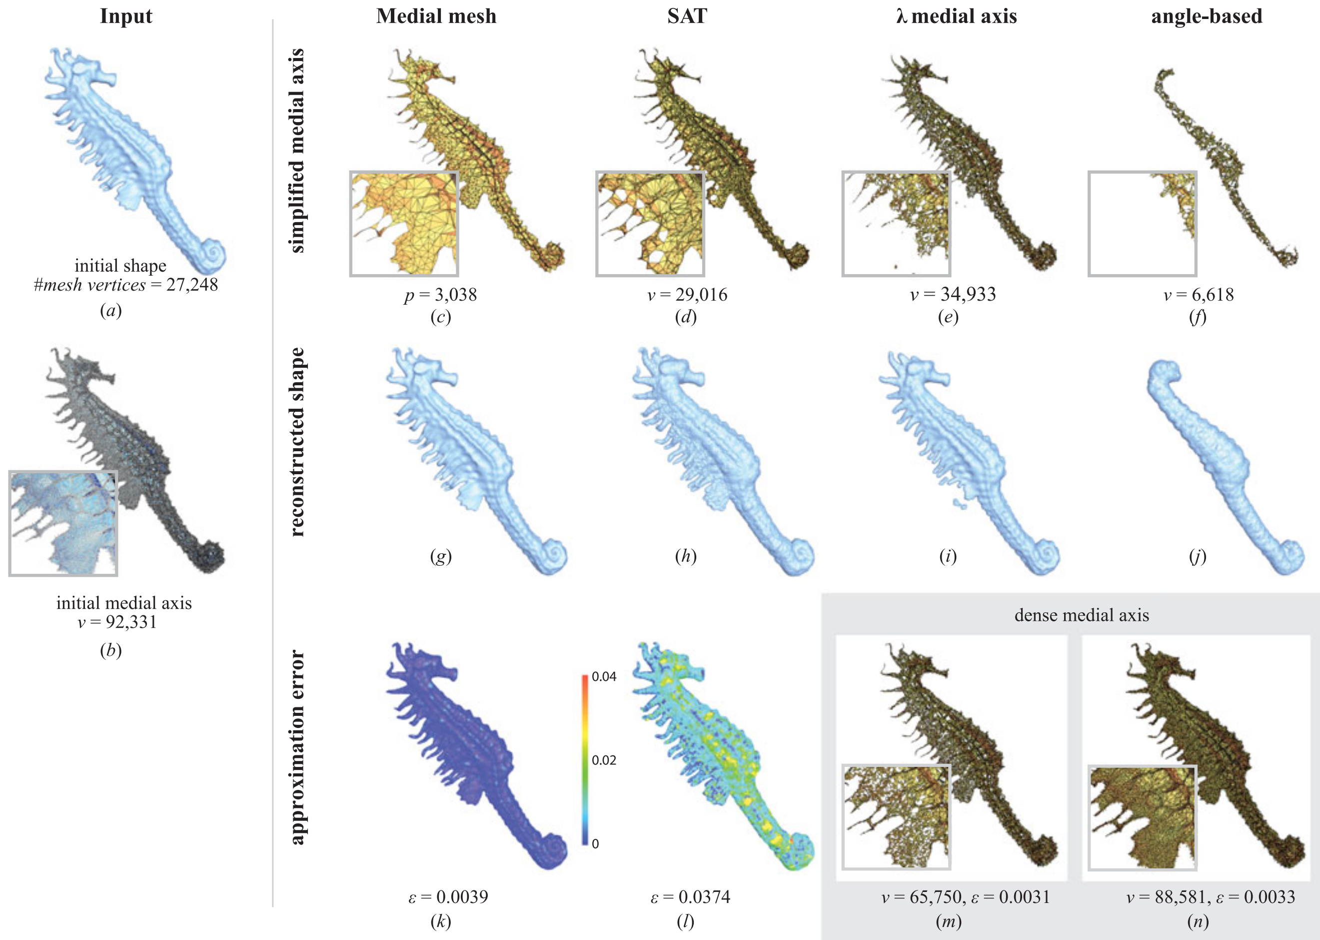

Fig. 7. Comparisons of 3D medial axis simplification methods. The first column (a and b) shows the original seahorse model and the initial Voronoi-based medial axis. The first row (c-f) shows the medial axes computed by the different methods. The second row (g-j) shows the boundary surfaces represented by the respective medial axes in the first row. (k) and (l) show the color-coded approximations of the surfaces in (g) and (h), respectively. (m) and (n) are the dense and unstable medial axes generated by the λ medial axis method and the angle-based method, respectively. (p: number of enveloping primitives, v: number of medial spheres, ε: approximation error.)

largely due to the different feature scales of the input shape (Fig. 6b). The angle-based method, on the other hand, does not preserve input topology and the resulting medial axis simply becomes disconnected (Fig. 6c). In addition, a main branch on the left has been incorrectly removed, therefore compromising significantly the approximation accuracy of the pruned medial axis. The medial mesh computed by our method, in contrast, uses much fewer sample circles while achieving the smallest approximation error among all. Visually, the union of envelopes of adjacent medial circles in the medial mesh yields a more accurate shape approximation, as shown in Fig. 6d. In comparison, gaps between the union of medial circles and the original shape boundary by the other methods are clearly visible. We also tested a variant of the angle-based method by enforcing topology preservation during filtering. It uses 2,978 medial points to achieve an approximation error of 0.01, while a medial mesh uses only 79 primitives for the same error level.

We next use a 3D shape to compare our method with the SAT method [11], in addition to the λ medial axis method and the angle-based method, as shown in Fig. 7. This 3D seahorse model has 27K vertices and its initial medial axis has many unstable branches (Fig. 7b). Our method significantly simplifies the medial axis to reduce the number of primitives to $\sim$3K; still this highly simplified medial mesh achieves an accurate shape reconstruction with an approximation error smaller than 0.004 only (Figs. 7c, 7g, and 7k). In contrast, the shape approximation error of the filtered medial axis from SAT is 0.037 even with $\sim$30K spheres, about an order of magnitude more primitives than our method. Moreover, the topology of the medial axis has not been preserved (e.g., the holes created at the lower backbone of the seahorse in Figs. 7d, 7h, and 7l). Clearly, medial axis filtering with both λ medial axis and angle-based methods led to an unacceptable shape representation or topological change (Figs. 7e, 7i, 7f, and 7j). Meanwhile, the number of spheres used by either method is at least an order of magnitude larger than that in our method. To attain the same shape approximation errors as our result, a huge number of spheres would be necessary with both methods, with the corresponding medial axes being highly unstable (Figs. 7m and 7n).

Admittedly, these three methods (SAT, λ medial axis and angle-based methods) are mainly designed for removing medial axis instability, rather than specifically for accurate shape approximation. Nevertheless, our simplification method based on the medial mesh achieves both accurate shape approximation and a stable medial axis representation simultaneously. This is due to the fact that medial axis instabilities correspond to small

TABLE 1
Number of Primitives Used for Representing the 3D Seahorse Model, with All Four Methods Using the Enveloping Representation Like a Medial Mesh

	# medial vertices	# primitives	error
Medial mesh	352	2,586	0.02219
SAT	8,131	50,161	0.02316
λ medial axis	58,221	355,672	0.02293
angle-based	48,158	300,632	0.02453

perturbations on the shape boundary, the removal of which would not incur a significant shape approximation error and, therefore, they can be effectively eliminated with our error-driven simplification.

We also examine the simplification power of our method irrespective to the use of an enveloping representation. To this end, we compute simplified medial axes using the other three methods, and build the same enveloping representation from their resulting medial vertices just like the medial mesh. Table 1 shows that the medial mesh still uses the fewest number of primitives to attain the same approximation error among all four methods.

Although SAT may often result in topologically incorrect shape representation, an issue acknowledged in [11], it generally produces good shape approximation, though with a large number of vertices. In comparison, our method achieves the same approximation accuracy using far fewer primitives than SAT and preserves shape topology. This is important for many applications such as shadow computation and shape deformation where the computational complexity depends on the number of primitives. Table 2 shows a comparison of our method against SAT on several 3D models in terms of the number of primitives used against the approximation accuracy that can be achieved. It can be seen that for highly accurate approximation ($\varepsilon = 0.001$), SAT requires at least two orders of magnitude more primitives than a simplified medial mesh. Even for the less accurate approximation ($\varepsilon = 0.032$), the number of primitives needed by SAT is still more than one order of magnitude larger than the simplified medial mesh.

Comparison with feature-based methods We now compare our method with the feature-based medial axis

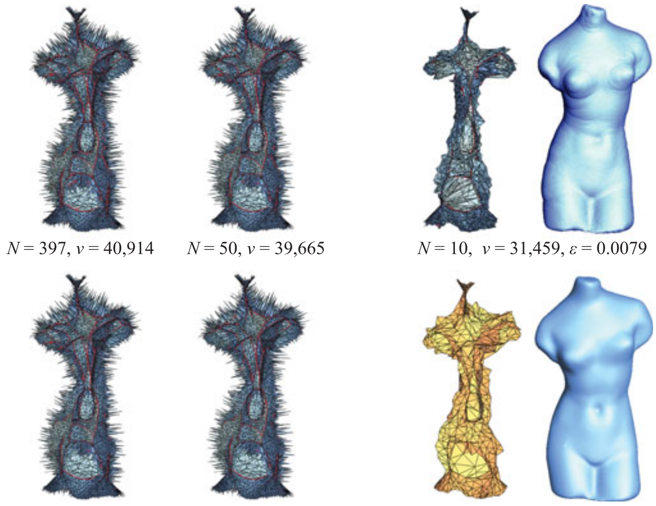
$N = 397$, $v = 40{,}914$ $N = 50$, $v = 39{,}665$ $N = 10$, $v = 31{,}459$, $\varepsilon = 0.0079$
$N = 26$, $v = 39{,}409$ $N = 10$, $v = 38{,}298$ $p = 1{,}735$, $\varepsilon = 0.0078$
(a) (b)

Fig. 8. (a) The feature-based method, Tam and Heidrich [2002], progressively reduces the number of medial sheets (N) to achieve shape simplification, taking a noisy Voronoi diagram of the shape as the initial input. The instability of the medial axis and the number of vertices (v) are not reduced significantly as the medial sheets are gradually removed. The non-manifold edges are marked in red. (b) Top: With a clean medial axis as the input, this method removes small features without reducing the number of medial vertices much. Bottom: An output achieving the same approximation error by our method is shown as a comparison.

simplification method by Tam and Heidrich [37] which prunes insignificant shape features based on a volume threshold. With a small volume threshold, the method is able to remove insignificant features while keeping relatively large sheets of medial axis representing important parts of the original shape. If the volume error threshold increases, significant pieces of the shape will begin to be trimmed off. It is worth noting that the method of Tam and Heidrich does not prune the unstable spikes and thus needs a relatively clean medial axis to start with [37]. Otherwise, if initialized with a highly unstable unfiltered medial axis such as that from the Voronoi diagram, this method is often unable to prune unstable branches that are not separate sheets. In any case, the remaining sheets are still a dense representation containing a large number of triangles (Fig. 8). Therefore, this method is not effective for

TABLE 2
Comparisons of the Approximation Error against the Number of Primitives in the Simplified Medial Axis for the SAT and Our Methods on Five 3D Models

		Approximation error					
		0.032	0.016	0.008	0.004	0.002	0.001
retinal	SAT	8,075	27,251	89,378	277,990	827,632	1,928,799
	Medial mesh	178	294	498	855	1,625	3,035
bird	SAT	6,785	25,365	98,771	383,880	1,392,083	5,086,712
	Medial mesh	265	478	839	1,479	2,296	5,503
table	SAT	12,833	52,463	202,814	773,646	2,971,080	11,542,309
	Medial mesh	116	327	886	2,378	5,645	10,759
girl	SAT	11625	37701	127557	430056	1276757	3245069
	Medial mesh	506	776	1476	2652	3913	7128
fandisk	SAT	19,344	71,818	273,847	1,047,419	3,875,252	13,130,365
	Medial mesh	242	398	587	904	1,742	3,059

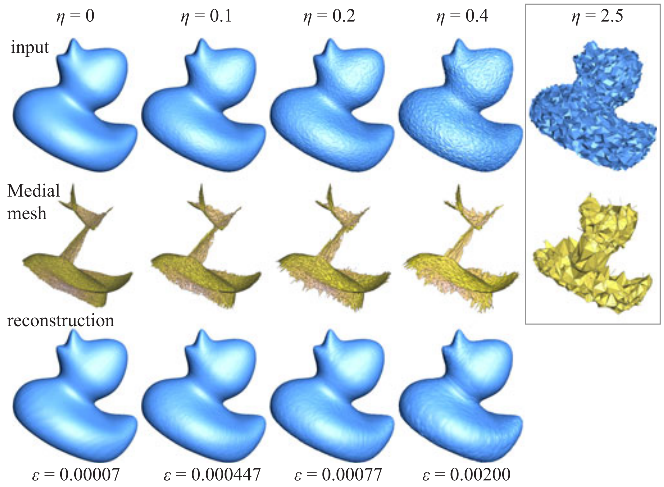

Fig. 9. Medial meshes for a noisy model on different levels of noise. The topology of the medial meshes is resistant to noisy input. Noises are removed effectively from the accurate shape representation. (ε is the Hausdorff distance between the surface represented by the simplified medial axis and the noise-free input.)

the purpose of geometric simplification; rather, it is primarily for simplifying the structure of a medial axis by removing those parts that are insignificant for representing the boundary. Our method, on the other hand, computes an accurate yet compact simplified representation.

Noisy models As MAT is in general sensitive to shapes with noisy boundaries, we next investigate how the medial mesh simplification is affected by noisy models. We apply different levels of noise relative to the average edge length of an input model (measured by $\eta \in [0, 1]$) and compute the corresponding simplified medial mesh. The result is shown in Fig. 9. The medial mesh simplification is capable of obtaining a stable medial representation and at the same time removing noise effectively as much as to a noise level of $\eta = 0.2$. The medial mesh retains very similar shape despite the increasing noise level. Also, it faithfully reproduces the original unnoisy input as indicated by the small distance error between the original shape and the shape represented by the simplified medial axis. With very noisy input, however, the simplified medial mesh is not a stable representation any more, as can be seen from the framed example in Fig 9. Our method fails in such case because the Hausdorff distance between the enveloping surface of the medial mesh and the highly unsmooth surface is no longer a reliable measure of the approximation error which is central to our simplification algorithm. Furthermore, since our algorithm preserves medial mesh homotopy, topological *noise* (for example, a small hole in a thin planar region) inherited by the initial medial mesh would persist and could not be removed.

More simplification results Fig. 1 shows that our method is capable of generating a series of simplified medial mesh at progressive levels of simplification with controlled approximation error. Note that the basic shape of an object is still preserved even when only a small number of medial spheres are kept. Fig. 10 shows more examples computed by our method. It can be seen that our simplification method can effectively generate compact and stable medial meshes that approximate shapes accurately within a user-specified error threshold.

6 Applications

We shall demonstrate the relevance of medial meshes in two shape modelling tasks. In shape approximation, the medial mesh is able to achieve much smaller approximation errors with the same number of primitives in comparison to state-of-the-art methods based on spherical approximation. In free-form shape deformation, the medial mesh can serve as a compact embedded structure of an object for direct manipulation to easily achieve thickness-preserving shape deformation.

6.1 Shape Approximation

The medial mesh proposed in this paper provides an effective alternative to shape approximation, compared to the conventional spheres representations [38], [39]. The key difference from the previous methods is that we use the interpolation of the medial spheres to approximate a given shape; in other words, the given shape is approximated with the union of the enveloping primitives defined by the medial mesh. Specifically, given any 3D shape to be approximated and an error tolerance provided by the user, we run our method for computing and simplifying the medial mesh. Once the simplified medial mesh is obtained, we collect all the enveloping primitives defined by the triangle faces of the mesh as well as those primitives defined by medial edges if the edges are not contained in any face. Then the union of these primitives is used as an approximation of the given shape, meeting the specified error tolerance.

Due to the use of simple sphere interpolation, the new approximation enabled by the simplified medial mesh we compute is much more efficient than the previous methods based on the union of spheres, as shown by the comparisons in Table 3 and Fig. 11. Typically, using the same number of primitives, the approximation error of the medial mesh is one order of magnitude smaller than those of the sphere-tree [39] and the medial spheres [38] methods, as shown in Table 3.

The enveloping primitives used in medial meshes are a new type of simple primitives capable of efficient and compact shape approximation, and have the potential to benefit important applications in graphics and geometry processing, such as fast shadow computation [40], [41]. While spheres are by themselves very simple primitives, further research would be needed of fast geometric computation involving the enveloping primitives used by the medial mesh.

6.2 Medial-Based Shape Deformation

The medial axis has been broadly adopted to provide a shape skeletal structure for shape deformation. While most methods use curve-based skeletons as the control structure (e.g., [28], [42]), others exploit the more general sheet-based medial representation (e.g. [27], [43]). Thanks to its explicit radius function, manipulation of medial axis can naturally give rise to a thickness preserving shape deformation, which is difficult to achieve under other free-form shape deformation paradigms. However, direct and coherent manipulation is non-trivial for complex medial structures. Special handling are needed for those two-sided medial

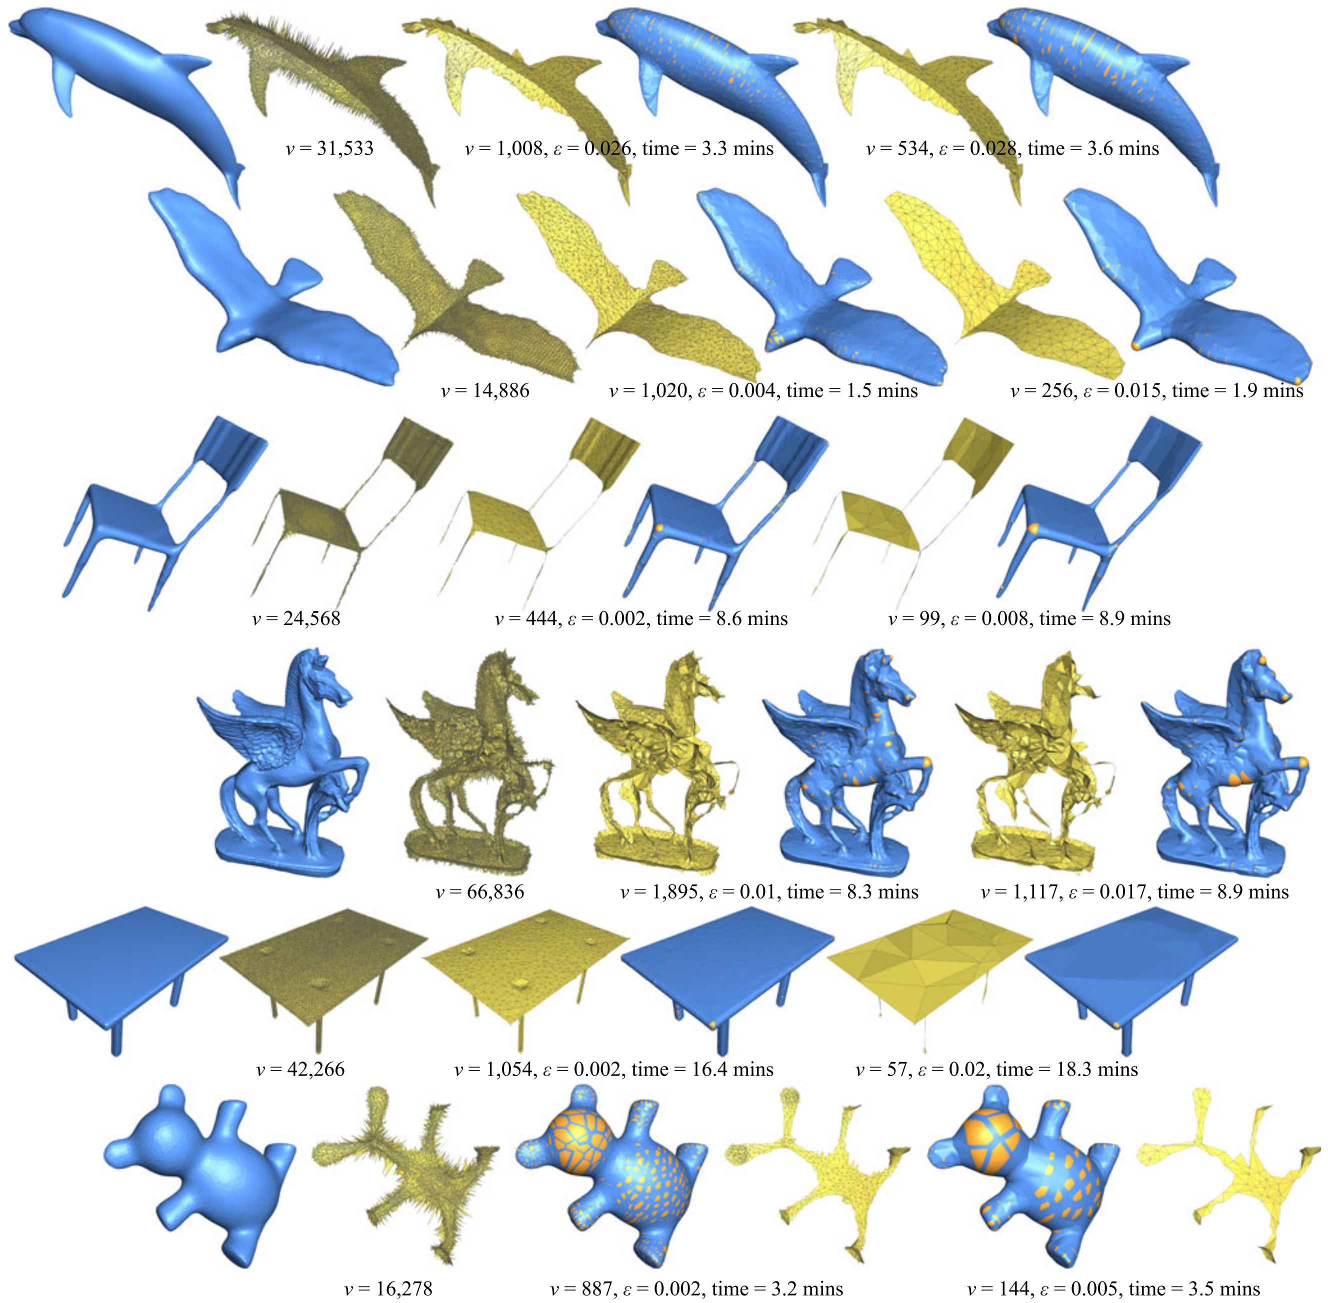

Fig. 10. Medial meshes computed by our simplification algorithm. The initial medial mesh and two levels of simplification with their corresponding number of medial spheres (v) and approximation error ε are shown for each shape.

axes computed by volume-collapsed methods [27], [30]). For instance, an additional stick-figure skeleton is required for consistent deformation of the two-sided skeleton medial mesh in [27].

Here, we use our medial mesh as an embedded graph that defines a subspace for a shape for embedded deformation proposed by Sumner et al. [44]. Changes to handles on a shape induces a deformation of the subspace, which results in the deformation of the shape itself. An advantage of this approach is that it provides intuitive direct control of the shape while leaving the embedded structure completely transparent to the users. Given a shape $\mathcal{S}$ as a mesh of vertices $\{\mathbf{v}_i\}$ and let $\mathcal{M}$ be its simplified medial mesh, whose vertices and edges constitute an embedded graph $\mathcal{G}$ of $\mathcal{S}$. Each medial vertex (or node) is associated with a local affine transformation. Similar to skeletal-subspace deformation (SSD), a vertex $\mathbf{v}_i$ on $\mathcal{S}$ is under the influence of the nodes in $\mathrm{Gov}_{\mathcal{M}}(\mathbf{v}_i)$ as well as their 1-ring neighboring nodes.

TABLE 3
Volume Difference with Respect to the Original Shape

	sphere-tree	medial spheres	Medial mesh
Duck	> 0.050	= 0.029	= 0.004
Venus	> 0.054	= 0.033	= 0.004
Bunny	> 0.091	= 0.062	= 0.007

There are 500 primitives used for all three methods and all three models. Data for the sphere-tree and medial spheres methods are taken from [38, Table 1]. The volume difference of our method is obtained by approximating each volume primitive with densely sampled spheres and computing the volumetric error to the original surface.

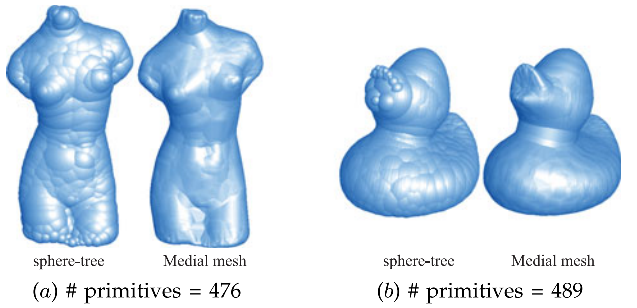

Fig. 11. Two models, (a) Venus, and (b) Duck, used for comparison of shape approximation in Table 3. The same number of primitives are used for both methods in each case.

The deformed position of $\mathbf{v}_i$ is a weighted average of the positions predicted by the transformations associated with these nodes, where the weights are assigned by considering the distances from $\mathbf{v}_i$ to the nodes. Following [44], we augment the edge set of $\mathcal{G}$ with edges connecting nodes associated with a common mesh vertex. Since the medial vertices are only sparsely connected, we also include additional edges connecting the 2-ring neighboring vertices in the medial mesh in order to improve transformation consistency among nearby nodes.

The deformation of the embedded graph is done via a minimization framework that considers as-rigid-as possible local transform ($\mathrm{E_{rot}}$), regularization for consistency between neighboring transforms ($\mathrm{E_{reg}}$), and the positional constraints imposed by the handles ($\mathrm{E_{con}}$) [44]. We introduce an extra energy term ($\mathrm{E_{twi}}$) for regularizing twisting along dangling edges representing tubular regions of $\mathcal{S}$, which requires that the difference between the end rotations of a dangling edge be minimal. This is necessary since the local transform along a dangling edge is ambiguous as the rotation around the edge is unspecified.

Vertex positions of the deformed mesh are obtained by vertex blending as in SSD, which in general does not respect the radius function of $\mathcal{M}$. A simple treatment is applied to achieve thickness preservation of the deformed shape. An offset distance to $\partial E(\mathcal{M})$, the enveloping surface of $\mathcal{M}$, is kept for each vertex. Deformed vertices are projected on the deformed enveloping surface and the offset distance is restored.

We compare our method with the embedded deformation originally suggested in [44] which uses evenly distributed samples drawn on $\mathcal{S}$ to form an embedded graph (which we named surface graph). Each mesh vertex is associated with its four-nearest nodes in the graph, and edges of the graph connect every two nodes that have a common association with a mesh vertex. We also compare against the CGAL implementation of the as-rigid-as possible method by Sorkine and Alexa [45] which aims to preserve local shape by minimizing the deviation from rigid transformation defined at each mesh vertices over its adjacent triangles. Fig. 12 shows the deformation results. Our method using medial mesh as an embedded graph leads to deformation that generally preserves thickness of the objects. By fixing the base of the model and moving only the handle at the top of the fertility model sideways, a nice bending effect can be obtained naturally. For surface graph embedded deformation, since graph nodes at the back and on the tummy of the dolphin are not connected, the surface at these two regions is easily pulled apart and the body is enlarged. An undesirable global shear in the fertility model is also observed and more sophisticated handle manipulation is needed to alleviate the effect. The surface graph also preserves well the thickness of the bird wing since the embedded graph nodes on the two sheets of the wings are closely connected. While the as-rigid-as-possible method can preserve local rigidity as much as possible, there is obvious change in the thickness, a non-local property, in the dolphin example, and it also results in the intersection of the two sheets of the wing in the bird example as well.

We note here that the deformation results of these methods depend also on the handle control. The surface graph embedded deformation and the as-rigid-as-possible method undoubtedly have their own merits in achieving high quality free-form deformation. By these examples we aim to demonstrate the effective use of the medial mesh as a medial structure for thickness-preserving deformation. However, many application-specific deformation requirements such as surface details preservation, etc., are yet to be fulfilled.

As a final remark, an unoptimized implementation of our method takes about 1 second for the bird deformation (with 20K mesh vertices, 220 medial vertices, 427 medial edges plus 294 augmented edges in the embedded graph), in which about 60 percent of the time is used in optimizing the embedded graph, and the remaining 40 percent is for obtaining deformed mesh vertex positions. This is largely due to the dense graph connection among the medial vertices, as well as the heavy node association for each mesh vertex. More investigation is also needed to further improve the time performance.

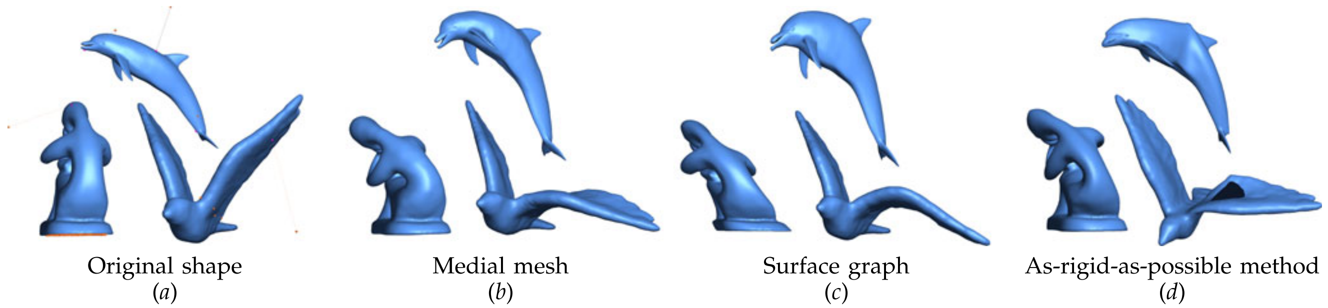

Fig. 12. Result of different deformation methods. (a) Point handles (red) define the user-specified constraints with the intended deformed positions (orange). Embedded deformation using (b) medial mesh and (c) surface graph as embedded graph; and (d) as-rigid-as-possible shape deformation.

7 CONCLUSION

We have proposed the medial mesh as a new discrete approximation of the medial axis. The medial mesh defines a compact representation of a 3D shape as the union of simple enveloping primitives generated by swept spheres. We have also presented an efficient algorithm for computing a simplified and stable medial mesh of a given 3D shape. Experiments show that our method is efficient and robust. The medial meshes computed by our methods are much simpler and offer more accurate shape approximation than the results by previous methods. We have presented applications of the medial mesh to shape approximation and shape deformation. For shape approximation, the medial mesh is shown to provide a much better approximation than the existing methods using the union of spheres. For shape deformation, due to its simplicity and intrinsic nature, the medial mesh demonstrates better performance in shape thickness preservation.

In summary, the medial mesh is a compact and stable representation of the medial axis, and thus has overcome the two notorious drawbacks of the medial axis, namely, instability and redundancy. Given the importance of the medial axis as a powerful intrinsic shape descriptor, we believe that the medial mesh will find more applications in shape modelling and analysis.

Further improvement of the medial mesh is possible. The enveloping surface of the medial mesh is only G^0 continuous. Therefore, one future problem is to use piecewise smooth surface (such as subdivision surfaces) to approximate the medial axis to achieve higher order and smoother shape approximation. Another potential improvement of the medial mesh is its mesh connectivity. Our method for simplifying a medial mesh resembles the paradigm of mesh simplification based on edge merging. The resulting medial mesh, while stable and simple, may be further improved by optimizing its mesh connectivity or mesh vertex distribution, similar to the effect of surface remeshing.

APPENDIX

POINT TO ENVELOPING PRIMITIVE DISTANCE COMPUTATION

Let $\mathcal{M} = \{m_i\}$ be a medial mesh approximating a shape $\mathcal{S} \in \mathbb{R}^3$, where m_i denote a simplex in $\mathcal{M}$. We now consider the distance of a point $\mathbf{q} \in \partial\mathcal{S}$ to an enveloping primitive $E(m_i)$. Let us assume that $m_i = (\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)$ is a medial face, the most general element in $\mathcal{M}$; the cases of m_i being a medial vertex or edge then follow similarly. The surface $\partial E(m_i)$ consists of patches from three spheres, three truncated cones and two triangle faces, one for the hull on each side of m_i (Fig. 2 right). The equations of the truncated cones and triangles can be derived and the signed distance to them can be computed. We therefore have:

1) three signed distances $d^*(\mathbf{q}, \mathbf{v}_j)$ to the spheres at $\mathbf{v}_j$, $j = 1, 2, 3$;
2) three signed distances $d^*(\mathbf{q}, \mathbf{v}_j\mathbf{v}_k)$ to the truncated cones defined by edge $\{\mathbf{v}_j, \mathbf{v}_k\}$, $j, k \in \{1, 2, 3\}$;
3) one signed distance $d^*(\mathbf{q}, \mathbf{v}_1\mathbf{v}_2\mathbf{v}_3)$ to the triangle face on the hull lying in the same halfspace of the medial face $\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ as $\mathbf{q}$.

Note that in cases (2) and (3), if the footpoint of $\mathbf{q}$ does not lie within the extent of the truncated cone or the triangle, the signed distance will be equal to ∞.

The distance of $\mathbf{q}$ to $E(m_i)$ is then given by

$$d\big(\mathbf{q}, \partial E(m_i)\big) = |\min\{d^*(\mathbf{q}, \mathbf{v}_j), d^*(\mathbf{q}, \mathbf{v}_j\mathbf{v}_k), d^*(\mathbf{q}, \mathbf{v}_1\mathbf{v}_2\mathbf{v}_3) \mid j, k \in \{1, 2, 3\}\}|.$$

ACKNOWLEDGMENTS

The work of Wenping Wang was partially supported by the National Basic Research Program of China (2011CB302400), the National Natural Science Foundation of China (61272019, 61332015, 61202148), and the Research Grant Council of Hong Kong (718311 and 717813).

REFERENCES

[1] H. Blum, "A transformation for extracting new descriptors of shape," in *Models for the Perception of Speech and Visual Form*, W. Wathen-Dunn, Ed. Cambridge: MIT Press, 1967, pp. 362–380.

[2] M. Kazhdan, B. Chazelle, D. Dobkin, T. Funkhouser, and S. Rusinkiewicz, "A reflective symmetry descriptor for 3D models," *Algorithmica*, vol. 38, no. 1, pp. 201–225, Oct. 2003.

[3] K. Siddiqi, J. Zhang, D. Macrini, A. Shokoufandeh, S. Bouix, and S. Dickinson, "Retrieving articulated 3-D models using medial surfaces," *Mach. Vision Appl.*, vol. 19, no. 4, pp. 261–275, May 2008.

[4] K. Siddiqi and S. Pizer, *Medial Representations: Mathematics, Algorithms and Applications*. New York, NY, USA: Springer, 2008.

[5] T. Culver, J. Keyser, and D. Manocha, "Exact computation of the medial axis of a polyhedron," *Comput. Aided Geom. Des.*, vol. 21, no. 1, pp. 65–98, Jan. 2004.

[6] O. Aichholzer, W. Aigner, F. Aurenhammer, and B. Jüttler, "Exact medial axis computation for triangulated solids with respect to piecewise linear metrics," in *Proc. 7th Int. Conf. Curves Surfaces*, 2010, pp. 1–27.

[7] D. Attali and A. Montanvert, "Modeling noise for a better simplification of skeletons," in *Proc. Int. Conf. Image Process.*, Sep 1996, vol. 3, pp. 13–16.

[8] N. Amenta, S. Choi, and R. K. Kolluri, "The power crust," in *Proc. 6th ACM Symp. Solid Model. Appl.*, 2001, pp. 249–266.

[9] T. K. Dey and W. Zhao, "Approximate medial axis as a Voronoi subcomplex," in *Proc. 7th ACM Symp. Solid Model. Appl.*, 2002, pp. 356–366.

[10] M. Foskey, M. C. Lin, and D. Manocha, "Efficient computation of a simplified medial axis," in *Proc. 8th ACM Symp. Solid Model. Appl.*, 2003, pp. 96–107.

[11] B. Miklos, J. Giesen, and M. Pauly, "Discrete scale axis representations for 3D geometry," *ACM Trans. Graph.*, vol. 29, pp. 101:1–101:10, Jul. 2010.

[12] F. Chazal and A. Lieutier, "The λ-medial axis," *Graph. Models*, vol. 67, no. 4, pp. 304–331, Jul. 2005.

[13] J. Chaussard, M. Couprie, and H. Talbot, "A discrete λ-medial axis," in *Proc. 15th IAPR Int. Conf. Discrete Geometry Comput. Imagery*, 2009, pp. 421–433.

[14] N. Faraj, J.-M. Thiery, and T. Boubekeur, "Progressive medial axis filtration," in *Proc. SIGGRAPH Asia 2013 Techn. Briefs*, 2013, pp. 3:1–3:4.

[15] R. Tam and W. Heidrich, "Shape simplification based on the medial axis transform," in *Proc. IEEE Vis.*, 2003, pp. 481–488.

[16] D. Attali and J.-O. Lachaud, "Delaunay conforming iso-surface, skeletonextraction and noise removal," *Comput. Geometry*, vol. 19, no. 23, pp. 175–189, 2001.

[17] J. Bloomenthal and C. Lim, "Skeletal methods of shape manipulation," in *Proc. Int. Conf. Shape Model. Appl.*, 1999, pp. 44–47.

[18] E. Larsen, S. Gottschalk, M. C. Lin, and D. Manocha, "Fast proximity queries with swept sphere volumes," Univ. North Carolina at Chapel Hill, Chapel Hill, NC, USA, Tech. Rep. TR99-018, 1999.

[19] A. Guéziec, "'meshsweeper': Dynamic point-to-polygonal-mesh distance and applications," *IEEE Trans. Vis. Comput. Graph.*, vol. 7, no. 1, pp. 47–61, Jan. 2001.

[20] S. Yoshizawa, A. G. Belyaev, and H.-P. Seidel, "Free-form skeleton-driven mesh deformations," in *Proc. 8th ACM Symp. Solid Model. Appl.*, 2003, pp. 247–253.

[21] B. Bastl, B. Jüttler, J. Kosinka, and M. Lávička, "Volumes with piecewise quadratic medial surface transforms: Computation of boundaries and trimmed offsets," *Comput. Aided Des.*, vol. 42, no. 6, pp. 571–579, Jun. 2010.

[22] S. M. Pizer, D. S. Fritsch, P. A. Yushkevich, V. E. Johnson, and E. L. Chaney, "Segmentation, registration, and measurement of shape variation via image object shape," *IEEE Trans. Med. Imag.*, vol. 18, no. 10, pp. 851–865, Oct. 1999.

[23] O. Aichholzer, F. Aurenhammer, D. Alberts, and B. Grtner, "A novel type of skeleton for polygons," *J. Universal Comput. Sci.*, vol. 1, pp. 752–761, 1996.

[24] M. Tanase and R. C. Veltkamp, "Polygon decomposition based on the straight line skeleton," in *Proc. 19th Annu. Symp. Comput. Geometry*, 2003, pp. 58–67.

[25] E. Demaine, M. Demaine, and A. Lubiw, "Folding and cutting paper," in *Proc. Japanese Conf. Discr. Comput. Geometry*, 2000, pp. 104–118.

[26] G. Barequet, M. T. Goodrich, A. Levi-Steiner, and D. Steiner, "Straight-skeleton based contour interpolation," in *Proc. 14th Annu. ACM-SIAM Symp. Discrete Algorithms*, 2003, pp. 119–127.

[27] S. Yoshizawa, A. G. Belyaev, and H.-P. Seidel, "Skeleton-based variational mesh deformations," *Comput. Graph. Forum*, vol. 26, no. 3, pp. 255–264, 2007.

[28] O. K.-C. Au, C.-L. Tai, H.-K. Chu, D. Cohen-Or, and T.-Y. Lee, "Skeleton extraction by mesh contraction," *ACM Trans. Graph.*, vol. 27, no. 3, pp. 44:1–44:10, Aug. 2008.

[29] L. Liu, E. W. Chambers, D. Letscher, and T. Ju, "A simple and robust thinning algorithm on cell complexes," *Comput. Graph. Forum*, vol. 29, no. 7, pp. 2253–2260, 2010.

[30] J.-M. Thiery, E. Guy, and T. Boubekeur, "Sphere-meshes: Shape approximation using spherical quadric error metrics," *ACM Trans. Graph.*, vol. 32, no. 6, pp. 178:1–178:12, Nov. 2013.

[31] S. W. Choi and H.-P. Seidel, "Hyperbolic Hausdorff distance for medial axis transform," *Graphical Models*, vol. 63, no. 5, pp. 369–384, 2001.

[32] N. Amenta and M. Bern, "Surface reconstruction by Voronoi filtering," in *Proc. 14th Annu. Symp. Comput. Geometry*, 1998, pp. 39–48.

[33] J. Brandt, "Convergence and continuity criteria for discrete approximations of the continuous planar skeleton," *CVGIP: Image Understanding*, vol. 59, no. 1, pp. 116–124, 1994.

[34] T. K. Dey, H. Edelsbrunner, S. Guha, and D. V. Nekhayev, "Topology preserving edge contraction," *Publ. Inst. Math.*, vol. 66, pp. 23–45, 1998.

[35] S. Pion and M. Teillaud, "3D triangulations," in *CGAL User and Reference Manual*, 4.4 ed. CGAL Editorial Board, 2000.

[36] N. Kruithof, "3D skin surface meshing," in *CGAL User and Reference Manual*, 4.4 ed. CGAL Editorial Board, 2000.

[37] R. C. Tam and W. Heidrich, "Feature-preserving medial axis noise removal," in *Proc. 7th Eur. Conf. Comput. Vis.-Part II*, 2002, pp. 672–686.

[38] S. Stolpner, P. Kry, and K. Siddiqi, "Medial spheres for shape approximation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 6, pp. 1234–1240, Jun. 2012.

[39] G. Bradshaw and C. O'Sullivan, "Adaptive medial-axis approximation for sphere-tree construction," *ACM Trans. Graph.*, vol. 23, no. 1, pp. 1–26, Jan. 2004.

[40] Z. Ren, R. Wang, J. Snyder, K. Zhou, X. Liu, B. Sun, P.-P. Sloan, H. Bao, Q. Peng, and B. Guo, "Real-time soft shadows in dynamic scenes using spherical harmonic exponentiation," *ACM Trans. Graph.*, vol. 25, no. 3, pp. 977–986, Jul. 2006.

[41] R. Wang, K. Zhou, J. Snyder, X. Liu, H. Bao, Q. Peng, and B. Guo, "Variational sphere set approximation for solid objects," *Vis. Comput.*, vol. 22, no. 9, pp. 612–621, Sep. 2006.

[42] I. Baran and J. Popović, "Automatic rigging and animation of 3d characters," *ACM Trans. Graph.*, vol. 26, no. 3, pp. 72:1–72:8, Jul. 2007.

[43] J. Bloomenthal, "Medial-based vertex deformation," in *Proc. ACM SIGGRAPH/Eurograph. Symp. Comput. Animation*, 2002, pp. 147–151.

[44] R. W. Sumner, J. Schmid, and M. Pauly, "Embedded deformation for shape manipulation," *ACM Trans. Graph.*, vol. 26, no. 3, pp. 80:1–80:7, Jul. 2007.

[45] O. Sorkine and M. Alexa, "As-rigid-as-possible surface modeling," in *Proc. 5th Eurograph. Symp. Geometry Process.*, 2007, pp. 109–116.

Feng Sun received the BEng and MEng degrees from Shandong University and the PhD degree from The University of Hong Kong, all in computer science. His research interests include mesh generation, computational geometry, and computational topology.

Yi-King Choi received the BSc and MPhil degrees in computer science from the University of Hong Kong, Hong Kong, in 1996 and 2000, respectively. She is currently a research assistant professor in the Department of Computer Science, The University of Hong Kong. Her research interests include computational geometry, geometric computing, and visualization.

Yizhou Yu received the PhD degree from the University of California at Berkeley in 2000. He is currently a full professor in the Department of Computer Science, The University of Hong Kong, and an adjunct professor at the University of Illinois, Urbana-Champaign. He received the 2002 National Science Foundation CAREER Award and the Best Paper Award at 2005 and 2011 ACM SIGGRAPH/EG Symposium on Computer Animation. He is on the editorial board of *Computer Graphics Forum* and *International Journal of Software and Informatics*. He is the program chair of Pacific Graphics 2009, Computer Animation and Social Agents 2012, and has served on the program committee of many leading international conferences, including SIGGRAPH, SIGGRAPH Asia, and International Conference on Computer Vision. His current research interests include data-driven methods for computer graphics and vision, digital geometry processing, video analytics, and biomedical data analysis. He is a member of the IEEE.

Wenping Wang is professor of computer science at The University of Hong Kong. His research interests include computer graphics, visualization, and geometric computing. He has recently focused on mesh generation and polyhedral surface modeling for architectural design. He is a journal associate editor of *Computer Aided Geometric Design*, *Computers & Graphics*, *Computer Graphics Forum*, *IEEE Computer Graphics and Applications*, and *IEEE Transactions on Visualization and Computer Graphics* (2008-2012). He is the program chair of Pacific Graphics 2003, ACM Symposium on Physical and Solid Modeling 2006, and International Conference on Shape Modeling 2009, and conference chair of Pacific Graphics 2012, ACM Symposium on Physical and Solid Modeling 2013, and SIGGRAPH Asia 2013. He is a member of the IEEE.