

Discrete Scale Axis Representations for 3D Geometry

Balint Miklos
ETH Zurich / EPFL

Joachim Giesen
Friedrich-Schiller-Universität Jena

Mark Pauly
EPFL

Abstract

This paper addresses the fundamental problem of computing stable medial representations of 3D shapes. We propose a *spatially adaptive* classification of geometric features that yields a robust algorithm for generating medial representations at different levels of abstraction. The recently introduced continuous scale axis transform serves as the mathematical foundation of our algorithm. We show how geometric and topological properties of the continuous setting carry over to discrete shape representations. Our method combines scaling operations of medial balls for geometric simplification with filtrations of the medial axis and provably good conversion steps to and from union of balls, to enable efficient processing of a wide variety shape representations including polygon meshes, 3D images, implicit surfaces, and point clouds. We demonstrate the robustness and versatility of our algorithm with an extensive validation on hundreds of shapes including complex geometries consisting of millions of triangles.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling;

Keywords: medial axis, stability, scale axis, geometry representations, shape analysis

1 Introduction

Geometric shape representations are at the core of algorithms in computer graphics and geometric modeling. Two main classes are prevalent today: *Explicit* representations store a set of primitives, such as points, edges, or polygons, to define a continuous boundary surface using appropriate interpolation schemes. Since the latter commonly make use of a (local) parameterization, these representations are also referred to as *parametric* surfaces. The second main class are *implicit* representations that model a surface as a level-set of some scalar function defined over the embedding domain. A popular variant are distance functions that represent the (signed) distance to the surface for each point of the ambient space.

In this paper we focus on a third alternative: *Medial* representations, which can be considered a hybrid of explicit and implicit representations. They explicitly store primitives (e.g., points, lines or polygons) deeply in the interior of the shape that carry distance information to its boundary. As a result, medial representations give direct access to both the shape interior and its boundary, enabling a variety of applications in shape analysis, recognition, segmentation, abstraction, or meshing. For a detailed discussion on applications, we refer the reader to [Pizer et al. 2008], where also examples are shown in geography, visual arts, motion planning, medicine, biology, perception, or crystallography.

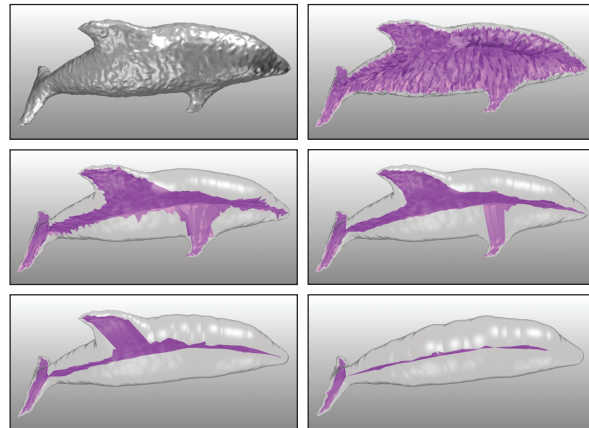


Figure 1: Different simplification levels of the discrete scale axis (purple) capture the dominant geometric features and allow shape representations at different levels of abstraction (transparent).

The most prominent medial representation is the medial axis transform [Blum 1967]. Given a shape $S \subset \mathbb{R}^d$, the medial axis M_S is the set of all points in the interior of S with at least two closest points on the boundary of S . For shapes in 3D, the medial axis is composed of 2D sheets that can degenerate into curves and single points. The medial axis transform also stores for every point on M_S its distance to the boundary of S . An important characteristic of this representation is bijectivity: Any closed shape has a unique medial axis transform and the boundary of a shape is defined exactly by its medial axis and the corresponding distance values. However, the medial axis lacks a crucial property, namely *stability*. Here stability means that two shapes that are geometrically close should also be close in terms of representation. Unfortunately, the medial axis is unstable: arbitrarily small changes of the shape can lead to large changes of its medial axis (see Figure 2). This renders the medial axis impractical in most applications, since geometric data is often imprecise due to noise or approximation in discrete representations.

A common approach to obtain practically useful medial representations is to discard parts of the medial axis that are deemed unstable, effectively trading exactness for stability. Typically, pruning is performed using a uniform local threshold at every medial axis point, e.g., on the angle formed by the vectors to the closest neighbors on the shape boundary [Attali and Montanvert 1996], or the circumradius of these closest neighbors [Chazal and Lieutier 2005]. However, it is well-known that this local filtering can lead to undesirable results for shapes with features at different scales [Pizer et al. 2003; Pizer et al. 2008; Attali et al. 2009] (see also Figure 2). We address this issue and propose a global approach that achieves significantly better results for a large class of shapes. Our method is based on the scale-adaptive ordering of features originating from the theory of the scale axis transform [Giesen et al. 2009], and defines a whole family of medial representations at different levels of abstraction, called *scale axis representations* (see Figure 1).

Contributions. The central contribution of this paper is a robust algorithm for computing stable medial representations that faithfully capture the geometric and topological features of a 3D shape. Our algorithm performs a multiplicative scaling of medial balls to

obtain a scale-adaptive ordering of geometric features. This construction, in conjunction with topology-preserving medial axis filtration techniques, allows the computation of practical scale axis representations for a specific class of shapes, the union of a finite number of balls. To complete our pipeline, we present provably good conversion algorithms to and from union of balls. Combining these two main components, our method becomes a general conversion tool from polygonal meshes, 3D images, point sets, and implicit functions to medial representations. To demonstrate the robustness and versatility of our algorithm we present an extensive empirical validation on hundreds of shapes, illustrating potential benefits for geometry processing applications.

Related Work. Numerous geometry representations and shape descriptors related to medial or local symmetry structures have been proposed in the past, including curve skeletons (see [Cornea et al. 2007] for an extensive survey), chordal axes [Prasad 1997], symmetry sets [Bruce et al. 1985], or the planar reflective symmetry transform [Podolak et al. 2006]. We refer the reader to the recent book [Pizer et al. 2008] for a detailed review and focus here on approaches directly based on the medial axis. We restrict our discussion to algorithms that are not limited to 2D geometry, but suitable for general 3D shapes also.

Robustly computing the exact medial axis is feasible only for a limited class of 3D shapes. Examples are the algorithm of Amenta et al. [2001] to compute the medial axis of a union of balls and the method introduced by Culver et al. [2004], which can handle polyhedrons of up to a few hundred faces. More common are approximation algorithms that have been proposed for a wide variety of geometry representations. Siddiqi et al. [2002] compute medial axis approximations on 3D voxel grids by detecting shocks in a numerically tracked level set of the distance function using flux computations. Another strategy labels voxels as belonging to the medial axis based on functions evaluated on voxel neighborhoods and iterative voxel pruning (e.g. [Svensson 2001]). In computational geometry, various methods have been introduced that process a sampling of the shape boundary given as an unorganized point set [Amenta et al. 2001; Dey and Zhao 2004; Chazal and Lieutier 2005]. Typically, such algorithms output a carefully chosen subset of Voronoi faces or weighted Delaunay triangles derived from the input points. Polygonal representations are either converted to a sample set [Dey et al. 2003] or a voxel image [Sud et al. 2005; Foskey et al. 2003; Stolpner and Whitesides 2009], or processed directly as in [Belyaev et al. 2007]. However, the latter approach produces results that typically depend strongly on the meshing of the input shape.

Numerous constructions have been proposed to address the instability of the medial axis. The majority of these methods compute an importance value for every medial axis point and discard medial branches with importance below a given threshold. The λ -medial axis [Chazal and Lieutier 2005; Chaussard et al. 2009] uses the circumradius of the closest boundary points as the importance value for this pruning operation. While this criterion yields a provably stable subset of the medial axis for small enough values of the global parameter λ , this approach does not respect different scales apparent in the object. Another popular criterion is based on the *object angle*, i.e., the angle formed by a medial point with its closest boundary points [Attali and Montanvert 1996; Amenta et al. 2001; Dey and Zhao 2004; Foskey et al. 2003; Sud et al. 2005]. The object angle adapts better to the local geometric scale, but yields simplifications of different topological types (see Figure 2).

Both object angle and circumradius are local criteria in the sense that they consider geometric information extracted from a single medial ball to determine the importance of the corresponding medial axis point. In order to detect higher level saliency of me-

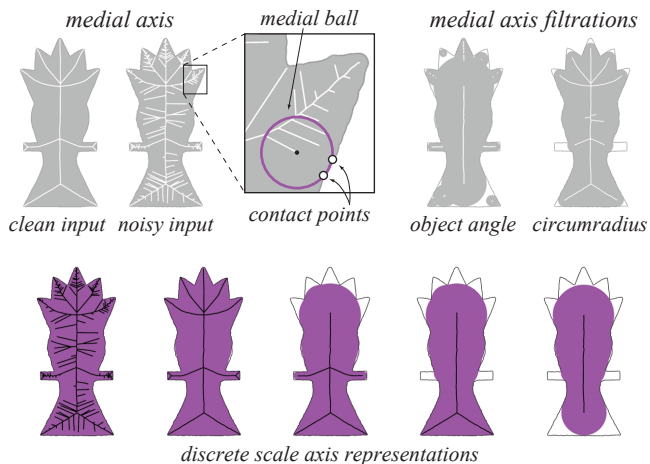


Figure 2: *Instability of the medial axis. Shapes that are geometrically close can have drastically different medial axes. Existing criteria to mitigate instability based on pruning by object angle and circumradius of the closest points produce unsatisfactory results if the shapes have features at different scales. The discrete scale axis representations locally adapts to the feature size, yielding more salient shape simplifications at different levels of abstraction.*

dial axis points, several global methods exploit the fact that a 2D shape is implicitly decomposed into two or more parts at every medial ball (see the Domain Decomposition Lemma in [Choi et al. 1997]). Unfortunately, such a property does not hold for 3D geometry (see Section 6.2.3 in [Pizer et al. 2008]) and the extension of the 2D methods is problematic. As an alternative, Tam and Heidrich [2003] compute saliency based on volume approximations, while Sud et al. [2005] use global connectivity information to control the topology of an object angle filtration. Ward and Hamarneh [2009] proposed statistical methods to identify stable medial points of an entire group of shapes, and Styner et al. [2003] exploit shape correspondences to compute stable and consistent representations. While these global 3D methods perform well in specific domains where multiple shapes of the same type are available (e.g. medical data sets of multiple patients), they are not applicable for the general problem of finding a stable medial representation of a single 3D shape that we consider here.

Instead of pruning the medial axis, another strategy is to directly modify the input shape and consider the medial axis of the modified shape as the stable representation. Since smoothing the boundary implicitly removes many undesired medial points, various numerical methods incorporate smoothing in the distance field estimation [Rumpf and Telea 2002]. However, choosing the right amount of smoothing can be difficult or impossible for complicated shapes with salient features at multiple scales. In general, as numerous researchers comment (e.g. [Pizer et al. 2003; Pizer et al. 2008; Attali et al. 2009]), computing medial representations that are stable with respect to perceptually small changes of the shape remains an unsolved problem in 3D.

Our solution to this problem builds on the theory of the scale axis transform that uses a scaling of medial balls to identify significant features of the shape [Giesen et al. 2009]. The simplicity of this construction makes the scale axis transform amenable to a rigorous theoretical analysis. Giesen and co-workers present a detailed mathematical study of the shape evolution under the scaling of medial balls in the continuous setting, i.e., for infinite sets of balls. Using the theory of semi-concave functions and non-smooth analysis, they prove that the scale axis preserves the topology of the input for simplification factors s smaller than a shape-dependent bound.




Figure 3: Our processing pipeline for computing discrete scale axis representations for 3D shapes. ‘MAT’ denotes the medial axis transform while ‘MAT*’ indicates the medial axis transform followed by a topology-preserving angle filtration.

In our work, we exploit the simplicity of the scale axis construction to implement an effective algorithm for computing stable medial representations for discrete 3D shapes.

Overview. Our method for computing medial representations fits into the general class of algorithms that follow the approximation paradigm described in [Attali et al. 2009]. The main strategy is to approximate a given shape S by a shape U whose medial axis M_U can be computed exactly, and to simplify M_U yielding a stable medial representation of S . Our choice for the approximating shape U is a union of balls, a decision motivated by several reasons: Firstly, any solid in \mathbb{R}^3 can be approximated to arbitrary precision by a finite union of balls. Secondly, the union of balls is one of the very few shape classes for which the medial axis can be computed exactly. Most importantly, for this class of shapes, we can design an efficient scale-adaptive medial axis simplification algorithm as a discrete version of the continuous scale axis. This algorithm is a combination of multiplicative scaling operations from the definition of the scale axis, and a topology-preserving angle filtration of the medial axis. More precisely, we grow a subset of the medial balls of the union of balls, compute the medial axis of the grown shape, apply an angle filtration on this medial axis, and shrink a carefully sampled set of medial balls back to the original scale. This simplified medial axis transform is the *discrete scale axis representation*. We can compute an explicit or implicit representation of the simplified shape from the set of shrunken balls. The complete processing pipeline of our method is shown in Figure 3.

2 Conversion to Union of Balls

The first step of our pipeline is the conversion of an input shape S into a finite union of balls U that approximates S within some user-specified tolerance δ , in our setting a distance value expressed relative to the diagonal of the bounding box of the input. For this conversion we propose an output-sensitive algorithm with mathematical guarantees on the quality of the approximation and bounds on the number of balls necessary. More precisely, our conversion algorithm guarantees that the computed union of balls is close to the input shape in Hausdorff distance and has the same homotopy type. Similar results have been presented for shapes represented by a dense boundary sampling, first in [Amenta and Kolluri 2000] and later in more detail in [Amenta et al. 2001]. Our approach extends these results to other shape representations like polygonal meshes, implicit surfaces and 3D images. The basic idea is to find a set of “large” balls – called polar balls – deep within the shape, using carefully chosen sample points on the shape boundary.

Polar balls. Our algorithm exploits properties of the polar balls of a sampling of the shape boundary. These balls are derived from the

Voronoi diagram of the sample points. Let us recall the following definitions: a *Voronoi vertex* is a point with four or more closest sample points. The corresponding *Voronoi ball* is the unique ball centered at the vertex that contains the closest sample points in its boundary. For every sample point there is an *inner pole*, namely the farthest Voronoi vertex in the intersection of the shape and the Voronoi cell of the sample point. Amenta et al. [2001] prove in the context of surface reconstruction that the inner poles converge to the medial axis of the shape as the density of the sampling increases. They also show that the union of polar balls, i.e., the Voronoi balls corresponding to inner poles, provides a good approximation of a shape. These results hold for surface samples of smooth surfaces as defined by the ϵ -sampling condition [Amenta and Bern 1998].

Sampling. To exploit the results of Amenta et al. [2001], we need to compute an appropriate sampling of the boundary surface of the input shape S . We use the method of [Boissonnat and Oudot 2005] that incrementally builds a set of sample points and an approximating surface mesh, until the approximation tolerance δ is met. The samples are generated using a Delaunay refinement technique that maintains a surface mesh as a specific subset of the Delaunay triangulation of the sample points, the Delaunay faces restricted to the boundary of the input shape. The beauty of this algorithm lies in the fact that the shape needs to be given only through an oracle that evaluates a single predicate: decide whether a given line segment intersects the shape boundary an odd number of times, and if so, compute one of the intersection points. Such an oracle can be implemented robustly and efficiently for both explicit and implicit representations using appropriate spatial data structures as detailed in [Boissonnat and Oudot 2005]. Therefore our algorithm can process 3D images, polygonal meshes, and level sets. Inputs described as point sets can be resampled, e.g., use moving least-squares approximation to evaluate the predicate. Alternatively, if the point sample is dense and noise-free we can directly use the power crust algorithm for the conversion to a union of balls.


Conversion Algorithm. Since our sampling comes with connectivity information, we can apply a much simpler and more robust algorithm for identifying the inner poles than the original method of [Amenta et al. 2001] that has been designed for unorganized point clouds. The sampling step extracts a subset of the Delaunay triangles that tightly approximates the boundary of the shape. Thus we can easily classify the Delaunay tetrahedra as inner and outer tetrahedra with respect to S . For any given point, in particular any Voronoi vertex, we can decide at the cost of a point location query in the Delaunay triangulation, if the point lies inside an inner tetrahedron. This easily allows identifying the farthest inner Voronoi vertex for every sample. If the boundary sampling is dense enough, this identification of inner poles is provably correct. In the appendix we provide bounds on the geometric closeness between the input

shape S and the approximating union of balls U . We also present an asymptotic upper bound on the number of balls generated by our algorithm, and prove homotopy equivalence of S and U . However, even if the conditions that are necessary to derive these theoretical guarantees are not satisfied, our algorithm produces good results in practice as illustrated in Figure 6. The statistics show that the runtime and the size of the output is mainly influenced by δ . As a result, the runtime is proportional to the size of the output (i.e. the algorithm is output-sensitive), which effectively allows trading computation time for precision using the tolerance δ .

3 Discrete Scale Axis

The conversion of a shape S into a union of balls U provides us with a shape representation that allows computing an exact medial axis. The essential question now becomes which parts of this medial axis, and hence which features on the boundary surface, should be considered stable. The key property of our approach is that we address this question in a scale-adaptive way: the importance of a surface feature is not determined in isolation based on purely local criteria, but depends on the relation of a feature to neighboring features, where the neighborhood size itself automatically adapts to the geometry. We achieve this spatially adaptive feature classification using a simple scaling of medial balls, i.e., balls centered on the medial axis that touch the boundary of the shape. We first review the main concepts of the continuous scale axis transform to illustrate the effect of the scaling operation. Then we discuss the main algorithmic components of our method in the discrete setting.

Scale Axis Transform. The scale axis transform has been introduced in [Giesen et al. 2009] as a generalization of the medial axis transform for solid shapes in \mathbb{R}^d . The s -scale axis transform of a shape S is defined by the following construction: (i) compute the medial axis of S yielding the (typically infinite) union of medial balls, (ii) scale the radius of each medial ball by the factor $s \geq 1$, (iii) recompute the medial axis of the scaled union of balls, and (iv) re-scale the new medial balls by the inverse factor $1/s$. The resulting set of balls is the s -scale axis transform of S and the ball centers define the corresponding s -scale axis. For $s = 1$ the scale axis transform is identical to the medial axis transform, while larger values of s lead to increasingly higher levels of simplification (see Figure 2). Let us illustrate the effect of the scaling operation on a simple 2D example:



We observe that medial balls are covered completely by larger nearby balls when scaling each ball’s radius. On the other hand, small balls without larger balls in their vicinity are not covered. This can be interpreted as a small feature next to a large feature being less important than a small feature that has no significantly larger features nearby. Recomputing the medial axis achieves the simplification and removes all covered balls, since these do not touch the boundary anymore and are thus no longer medial. To maintain the same global geometric scale, the final step of the construction shrinks the latter medial balls by the inverse of the original scaling factor. As the figure above illustrates, the scale-adaptive classification allows discarding unimportant branches of the medial axis, while keeping the features that are essential to capture the semantics of the shape, even if these features are small.




Figure 4: Scaling medial balls. The medial axis transform of the union of the two purple balls consists of infinitely many balls, some of which are indicated as gray circles. Growing all these balls is equivalent to growing only those balls centered at the vertices of the piecewise linear medial axis (a). When shrinking, however, we need to sample the medial axis to obtain a finite union of balls that approximates the shrunken shape (b).

There is an interesting connection between the scale axis and the object angle stability measure. The object angle at some medial axis point can be directly derived from the scaling factor at which the corresponding medial ball is covered by the neighboring medial balls during the multiplicative growing. The crucial difference in the definition of the scale axis is that there *all* medial balls are scaled simultaneously, thus a ball can be covered by some distant but large medial ball, as well. Consequently, stability is determined in a more global fashion. Intuitively, angle filtration determines importance based on local shape of the object, while the scale axis uses the relative sizes of features.

Medial Axis of a Union of Balls. Although the medial axis is geometrically unstable, we argue in the appendix that the exact medial axis M_U of the the union of balls U is well suited for approximating the stable subset of the medial axis M_S of S . To compute M_U , we use the method of Amenta and Kolluri [2001] and incorporate the extensions proposed by Tam and Heidrich [2003]. Details of the implementation are provided in Section 4.

Growing Medial Balls. The key element in the construction of the scale axis transform is the scaling of all medial balls by the simplification factor $s \geq 1$. However, an algorithmic solution to this scaling operation is not obvious, since in general the number of medial balls is infinite. We solve this problem by exploiting the specific structure of the medial axis M_U . Since U is a finite union of balls, M_U is piecewise linear and can be represented as a simplicial complex [Amenta and Kolluri 2001]. As illustrated in Figure 4 on a small example, the growing of a specific *finite* subset of medial balls creates *exactly the same shape* as growing the *infinite* set of all medial balls. It is easy to verify that this set consists of the medial balls corresponding to the vertices of the simplicial complex describing M_U . Note that this property only holds for the medial axis of a finite union of balls, which is exactly how we approximate the input shape S . The medial axis of the union of balls after scaling cannot contain the centers of all covered balls, hence these parts of the medial axis will be discarded. Effectively, M_U is simplified by removing unstable branches to yield a stable approximation of M_S .

Filtering. The scale axis construction has the property that branches of the medial axis can not only be pruned, but they can also move or get straightened out. This allows simplifying branches without breaking topology, even when aiming for drastic simplifications. However, the movement of branches also has a drawback: If the boundary of the shape contains spherical patches (or circular arcs in 2D), small spurious branches can appear (see Figure 5). Since the boundary of U consists by construction of such spherical




Figure 5: The limited resolution of the finite union of balls can lead to spurious branches. A local filtering step based on the object angle criterion removes these discretization artifacts.

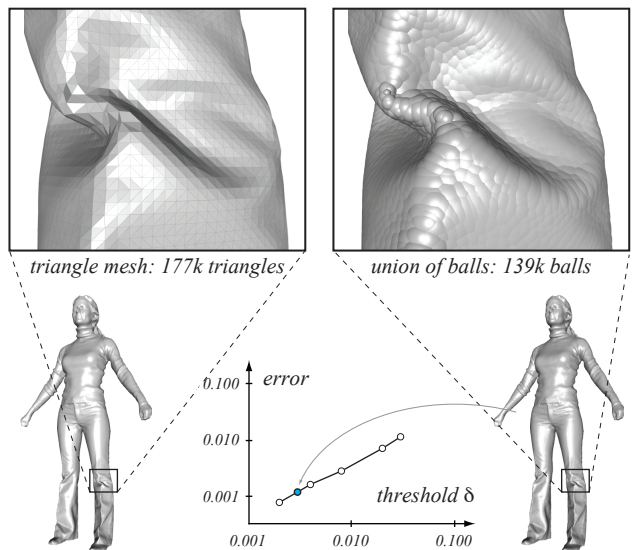
patches, we need to apply an explicit pruning step to remove these discretization artifacts. The pruning step is based on the observation that spurious branches appear when the intersection of a ball with the shape boundary consists not only of isolated points, but entire spherical patches. Any such ball needs to be medial and it may stay medial for a longer period when growing the radii, hence a new branch can originate at its location. However, the object angle of this medial ball decreases significantly, which allows identifying these spurious branches. To avoid changing the topology of the shape, we use a topology-conserving version of the angle filtration similar to [Sud et al. 2005]. As illustrated in Figure 5, this pruning step removes spurious branches without affecting the salient features of the scale axis representation.

Shrinking Medial Balls. The final step in the construction is re-scaling the medial balls of the grown shape by a factor $1/s$. Unfortunately, the finite scaling property described above only holds when growing balls, but not when shrinking them (Figure 4 (b)). In order to obtain a finite set of shrunken balls, we adaptively sample the scale axis such that the shrunken balls centered at these sample points intersect deeply, i.e., the intersection angle between two adjacent balls is above a certain threshold (we use 150° for all our examples). Edges of the simplicial representation of the scale axis are recursively split in half, while triangles are recursively subdivided using the standard 1-to-4 split, until balls centered at the triangle vertices meet at least as deeply as described by the threshold.

4 Implementation

Our implementation is based on CGAL (www.cgal.org). For the boundary sampling, we rely on the implementation of the method of Oudot et al. [2005] in the 3D meshing package of CGAL, and implement the polar ball extraction using the Delaunay triangulation package. The conversion from a union of balls to a polygonal mesh to reconstruct the simplified shapes is implemented using the 2D surface meshing package, treating the union of balls as an implicit surface and accelerating the point location using weighted Delaunay triangulations.

Medial Axis Computation. The largest part of the implementation is the method for computing the medial axis of union of balls as described in [Amenta and Kolluri 2001]. This algorithm combines computations on Delaunay triangulations and weighted alpha shapes, combinatorial tests involving different triangulations, and flooding algorithms to extract the medial axis. The main obstacle for a robust implementation are geometric degeneracies (multiple balls intersecting the same sample points). Tam and Heidrich [2003] have proposed various extensions to improve on robustness. For sets of polar balls they detect degenerate tetrahedra



δ	0.03	0.02	0.008	0.004	0.003	0.002
error	0.01137	0.00709	0.00281	0.00162	0.00119	0.00078
# balls	1,728	3,570	20,316	79,106	139,767	313,245
time (sec)	6.9	14.6	115.1	571.8	992.7	2,745.0

Figure 6: Statistics for the conversion from a polygon mesh to a union of balls and back to a mesh. The error is the Hausdorff distance between the input mesh and the output mesh relative to the diameter of the bounding box.

using only combinatorial information. We can apply this approach to compute the medial axis M_U of the union of polar balls U . Still, the geometric predicates on the input balls are degenerate by construction. While CGAL allows evaluating geometric predicates for degenerate configurations robustly, it does so at the cost of exact arithmetic computations. We reduce the time to evaluate such predicates considerably by perturbing the ball centers by a random displacement of at most a fixed factor of their radius (in our implementation we use a factor of 10^{-4}). Note that polar balls themselves are a sampling of the medial axis transform of U . Hence, one can improve efficiency without affecting the results by skipping the first medial axis computation and directly computing the medial axis of the grown set of polar balls.

Filtering. Since the computed medial axis has full connectivity information for its faces, our angle-based filtering step to remove spurious branches can be constrained to preserve topology, similar to the method of [Sud et al. 2005]. The algorithm first segments the medial axis into sheets that are bounded by non-manifold edges (edges where three or more medial faces meet) and assigns an importance value to every sheet. The importance value is the maximum object angle on this sheet. Sheets of minimal importance are removed iteratively, while maintaining the neighborhood information between sheets to preserve the topology of the medial axis. Since the segmentation of the medial axis in such sheets is unstable, this method is not suitable for general medial axis simplifications. In our setting, however, the spherical boundary patches induce only small spurious sheets that can be clearly separated and pruned without changing topology. Our implementations uses a fixed threshold on the object angle of 30° for the pruning.

Sheet boundaries. The approximation of S by a finite number of balls has a boundary composed of spherical patches, i.e., there are surface features at the intersections of adjacent balls resulting

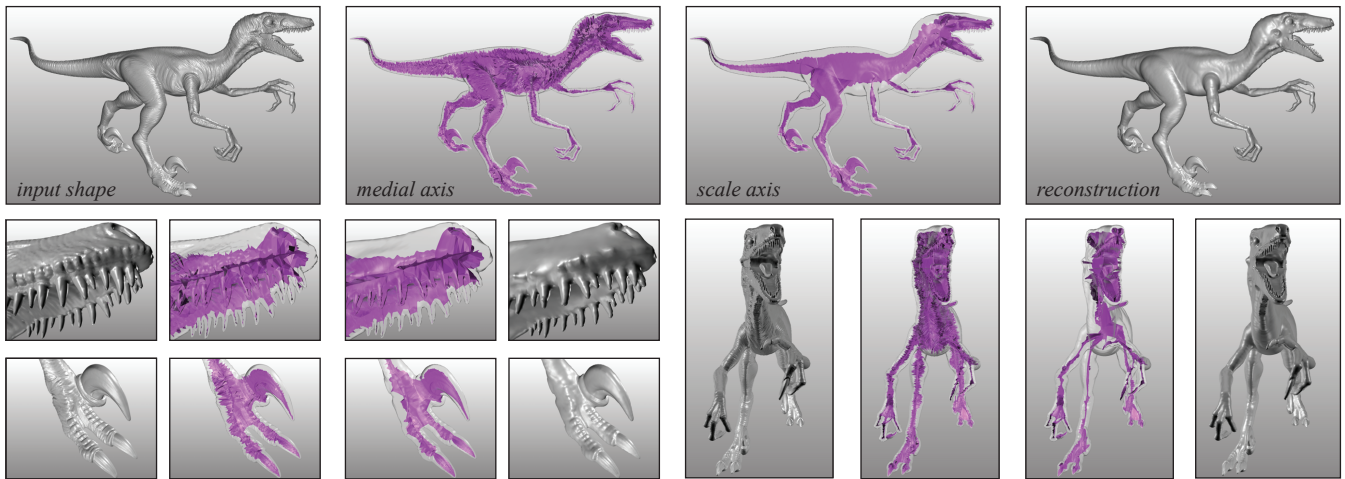


Figure 7: Comparison of the discrete scale axis representation for $s = 1.15$ and the original medial axis of the union of balls for a complex model consisting of 1.7 million triangles. Since unstable parts of the medial axis are removed in a scale-adaptive way, the representation preserves important features such as the teeth, while small wrinkles in the skin are removed.




Figure 8: A higher abstraction factor of $s = 1.2$ simplifies the medial axis more aggressively, removing most of the folds on the body but preserving the sharp edge of the pants, since this feature is dominant in relation to the geometry of its surroundings.

from the discretization that are not be present in the continuous medial axis transform. The scale axis representation captures both the salient features of the input shape and the ones resulting from the discretization. As a consequence, the sheets of the medial axis often have a jagged boundary depending on the distribution of the balls. For applications where more regular boundaries might be needed, we can smooth these boundaries by moving the centers of the medial balls that contribute to the sheet boundary and align them in a more regular fashion. Smoothing is implemented using the uniform Laplacian on the ball centers that lie on a one-manifold boundary of the medial axis. Since we ensure that the balls stay in the interior and touch the boundary of the unmodified union of balls, the shape is only altered minimally.

5 Evaluation and Discussion

In Figure 6 the performance of the sampling algorithm described in Section 2 is summarized. We show statistics of the conversion from

a mesh to a union-of-balls and back to a mesh for a complex shape with sharp features consisting of 177,044 triangles. As the zooms illustrate, geometric features are faithfully captured at a comparable number of primitives. The quality of the approximation, in addition to the theoretical guarantees and the robustness of the computation, make this conversion algorithm a useful tool in itself for various geometry processing applications.

Figures 7 and 8 show discrete scale axes for complex models with geometric features at different scales. The essential filtering step is the scaling of medial balls in conjunction with a re-computation of the medial axis to remove unstable branches in a scale-adaptive fashion. The simplicity of this construction is imperative for the effectiveness of the algorithm: no special cases need to be handled and no hidden parameters need to be adapted to specific models. Our method has two main parameters, the approximation threshold δ for the conversion to and from union of balls, and the scaling factor s that determines the level of simplification. Figure 9 shows how these parameters control the output of our algorithm. For high

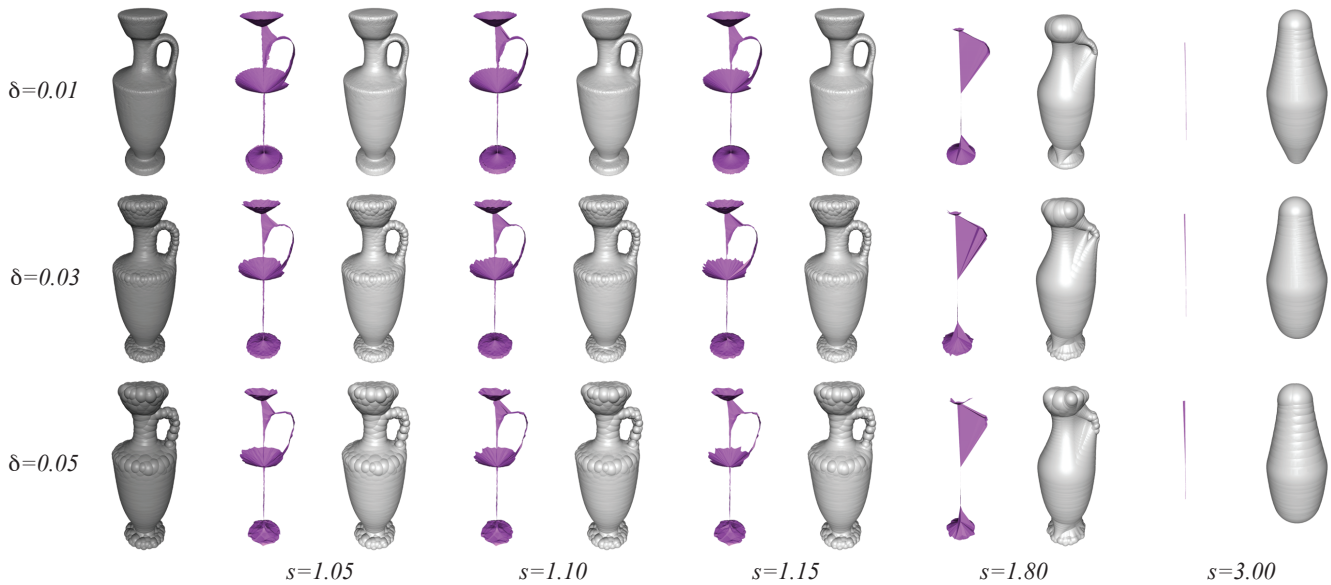


Figure 9: Medial representations for different values of the approximation threshold δ and scale parameter s .

levels of abstraction, i.e. large values of s , a higher threshold δ can be used for improved performance.

An important characteristic of the scale axis concept is that the resulting medial representation is stable for relatively long intervals of the scale parameter s , e.g., on Figure 9 the medial representation undergoes very minor changes for values of s in the interval $[1.05, 1.15]$. Hence, selecting suitable simplification levels becomes significantly easier than for existing methods that often exhibit a more volatile behavior when varying the filtering parameter (see video in additional material). Figure 11 provides a comparison with existing methods. Figure 12 summarizes results on a large set of different shapes ranging from mechanical parts to smooth organic surfaces. In total, we batch-processed all 380 models of the [Chen et al. 2009] database without any manual tuning of parameters or handpicking models. These examples confirm that the discrete scale axis is a general tool for geometric data analysis and processing.

Table 1 shows some statistics of our algorithm. Roughly 80% of the




Figure 10: Sharp features can be approximated up to the precision of δ and concave regions can be filled up for larger values of s .

Model (δ, s)	#Triangles	UoB	#Balls	DSA
Raptor (0.002, 1.15)	1,716,200	23.9	217,200	2.39
Woman (0.003, 1.2)	177,044	16.5	139,767	2.14
Shell (0.005, 1.1)	1,268,908	11.6	130,120	2.51
CAD (0.01, 1.2)	29,990	15.1	155,620	1.25
Vase (0.01, 1.1)	29,744	1.37	21,049	0.32
Vase (0.03, 1.1)	29,744	0.15	2,434	0.07
Vase (0.05, 1.1)	29,744	0.05	970	0.03

Table 1: Statistics and timings (in min.) measured on a Intel i7 2.93MHz with 6 Gb memory running our single-threaded 32-bit application. UoB denotes the conversion to a union of balls, while 'DSA' comprises the steps of the discrete scale axis computations.

time is spent on converting the input shape to a union of balls. Our current implementation focuses on robustness and is not optimized for speed, thus substantial performance improvements should be possible. For example, the sampling algorithm could be speed up by directly inserting the original vertices of the input mesh instead of exploring similar ones by the sampling algorithm.

Limitations. Figure 12 demonstrates that a wide range of models can be processed with our algorithm. In Figure 10 we illustrate some problems that can arise, in particular for large values of s and highly concave regions of the input shape boundary. These issues are related to the fact that the scale axis is not necessarily a subset of the exact medial axis and in extreme cases can even move outside the original shape. While this in general allows more drastic simplifications of the input shape, artifacts can appear in the form of unnatural merging of branches. A possible solution would be to explicitly map the discrete scale axis to a close subset of the medial axis and compute a medial axis filtration based on the scale axis. However, such a mapping between the scale axis and medial axis subsets is not straightforward to compute and needs further investigation.

While sharp features are handled robustly in practice as shown in Figure 10, an explicit representation of sharp features requires an

extension of our algorithm. Notice that the continuous scale axis transform does represent sharp features exactly, since the definition is based on an *infinite* number of balls. In the discrete setting, however, the finite number of balls is able to represent sharp features only for concave configurations. In future work we plan to investigate how to enrich a union of balls representation with additional primitives allowing to represent sharp features exactly, while still being able to efficiently compute the medial axis of the union.

6 Conclusion

We presented an effective method for generating discrete scale axis representations of complex 3D geometries. The conceptual simplicity of the central operation – scaling of medial balls – allows bridging the gap between sound mathematical theory and robust computations, yielding an effective algorithm suitable for a large class of shapes. Applications that so far have been hampered by a lack of robustness of such representations can benefit from our method. In particular, applications in shape analysis or geometric data processing that operate on noisy scans can employ the discrete scale axis representation, e.g. for feature extraction, segmentation, simplification, or abstraction. As a new general representation for 3D geometry, the discrete scale axis offers numerous avenues for future work. Interesting theoretical and practical questions remain to be investigated. For example, designing a variant of the scale axis as a nested sequence of geometric structures (instead of independent structures) would allow for directly extracting information encoded at all abstraction levels.

Acknowledgments We would like to thank to Camille Wormser for proof-reading the paper and to the AI Lab of University of Zurich for the Cardium Pseudolima shell model. This project has been supported by the Swiss National Science Foundation grant "The Scale Axis Transform" and DFG Priority Program 1335: Scalable Visual Analytics.

References

ALLIEZ, P., COHEN-STEINER, D., YVINEC, M., AND DESBRUN, M. 2005. Variational tetrahedral meshing. *ACM Trans. Graph.* 24, 3, 617–625.

AMENTA, N., AND BERN, M. 1998. Surface reconstruction by Voronoi filtering. *ACM Symposium on Computational Geometry.*

AMENTA, N., AND KOLLURI, R. K. 2000. Accurate and Efficient Unions of Balls. In *ACM Symposium on Computational Geometry*, 119–128.

AMENTA, N., AND KOLLURI, R. K. 2001. The medial axis of a union of balls. *Computational Geometry* 20, 1-2, 25–37.

AMENTA, N., CHOI, S., AND KOLLURI, R. K. 2001. The power crust, unions of balls, and the medial axis transform. *Computational Geometry* 19, 2-3, 127–153.

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

ATTALI, D., BOISSONNAT, J.-D., EDELSBRUNNER, H., AND BOISSONNAT, J.-D. 2009. Stability and Computation of the Medial Axis – a State-of-the-Art Report. *Math. Foundations of Scient. Visualization, Comp. Graphics, and Data Exploration.*

BELYAEV, A., YOSHIZAWA, S., AND SEIDEL, H.-P. 2007. Skeleton-based Variational Mesh Deformations. *Computer Graphics Forum* 26, 3, 255–264.

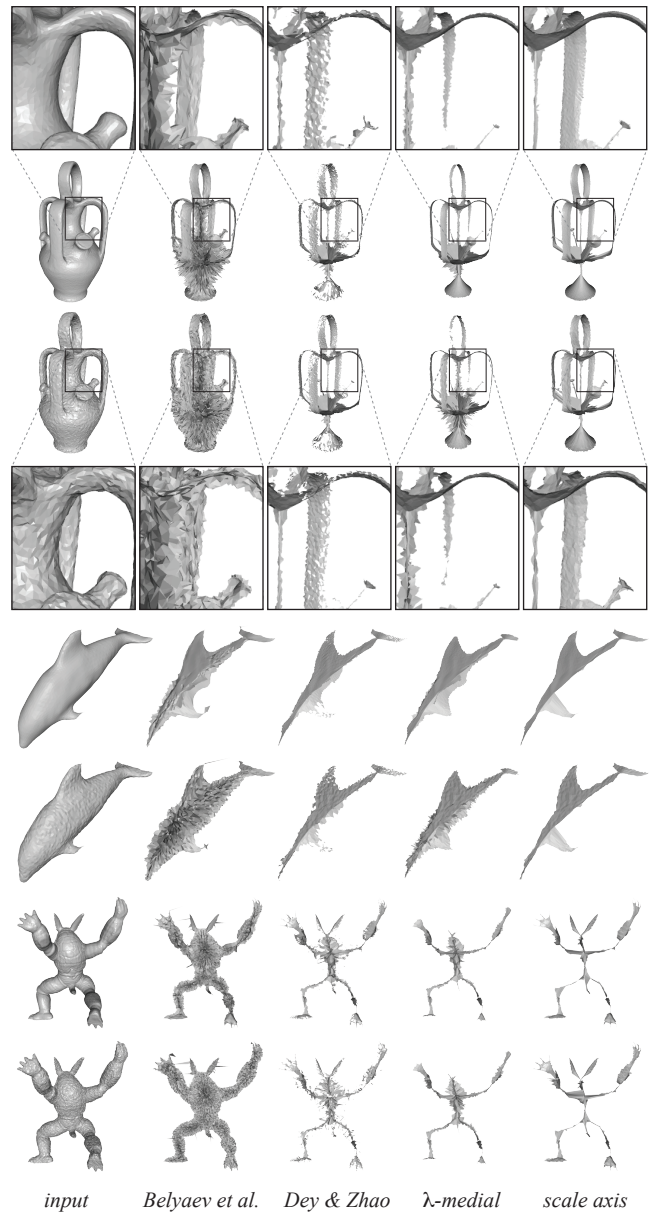


Figure 11: Comparison to other medial axis computation methods for different smooth and noisy models. For [Dey and Zhao 2004] and [Belyaev et al. 2007] we used the code provided on the authors’ website, the λ -medial axis is our implementation. Note that the uniform filtering of the λ -medial axis prunes away small, but important features like the tail of the Dolphin and ears of Armadillo, while the angle based method of [Dey and Zhao 2004] produces results of different topology than the input. If noise is not explicitly handled as in [Belyaev et al. 2007], results are clearly unsatisfactory.

BLUM, H. 1967. A Transformation for Extracting New Descriptors of Shape. *Models for the Perception of Speech and Visual Form*, 362 – 380.

BOISSONNAT, J.-D., AND OUDOT, S. 2005. Provably good sampling and meshing of surfaces. *Graph. Models* 67, 5, 405–451.

BRUCE, J. W., GIBLIN, P. J., AND GIBSON, C. 1985. Symmetry Sets. *Proc. of the Royal Society of Edinburgh* 101, A, 163–186.

CHAUSSARD, J., COUPRIE, M., AND TALBOT, H. 2009. A discrete lambda-medial axis. *DGCI* 5810, 421–433.

CHAZAL, F., AND LIEUTIER, A. 2005. The lambda-medial axis. *Graph. Models* 67, 304–331.

CHEN, X., GOLOVINSKIY, A., AND FUNKHOUSER, T. 2009. A benchmark for 3D mesh segmentation. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 28, 3 (Aug.).

CHOI, H. I., CHOI, S. W., AND MOON, H. P. 1997. Mathematical Theory Of Medial Axis Transform. *Pacific J. Math* 181, 57–88.

CORNEA, N. D., SILVER, D., AND MIN, P. 2007. Curve-Skeleton Properties, Applications, and Algorithms. *IEEE Transactions on Visualization and Computer Graphics* 13, 3, 18.

CULVER, T., KEYSER, J., AND MANOCHA, D. 2004. Exact computation of the medial axis of a polyhedron. *Computer Aided Geometric Design* 21, 1, 65–98.

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

DEY, T. K., WOO, H., AND ZHAO, W. 2003. Approximate medial axis for CAD models. *ACM Symposium on Solid and Physical Modeling*.

FOSKEY, M., LIN, M. C., AND MANOCHA, D. 2003. Efficient computation of a simplified medial axis. *ACM Symposium on Solid and Physical Modeling*.

GIESEN, J., MIKLOS, B., PAULY, M., AND WORMSER, C. 2009. The scale axis transform. *Annual Symposium on Computational Geometry*, 106–115.

PIZER, S. M., SIDDIQI, K., SZEKELY, G., DAMON, J. N., AND ZUCKER, S. W. 2003. Multiscale Medial Loci and Their Properties. *Int. J. Comput. Vision* 55, 155–179.

PIZER, S., SIDDIQI, K., AND YUSHKEVICH, P. 2008. *Medial Representations*, vol. 37 of *Computational Imaging and Vision*. Springer Netherlands, Dordrecht.

PODOLAK, J., SHILANE, P., GOLOVINSKIY, A., RUSINKIEWICZ, S., AND FUNKHOUSER, T. 2006. A planar-reflective symmetry transform for 3D shapes. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 25, 3, 549 – 559.

PRASAD, L. 1997. Morphological analysis of shapes. *CNLS Newsletter* 139, 1–18.

RUMPF, M., AND TELEA, A. 2002. A continuous skeletonization method based on level sets. In *VISSYM '02: Proceedings of the symposium on Data Visualisation 2002*, 151–ff.

SIDDIQI, K., BOUIX, S., TANNENBAUM, A., AND ZUCKER, S. W. 2002. Hamilton-Jacobi Skeletons. *International Journal of Computer Vision* 48, 3 (July), 215–231.

STOLPNER, S., AND WHITESIDES, S. 2009. Medial Axis Approximation with Bounded Error. In *International Symposium on Voronoi Diagrams in Science and Engineering*.

STYNER, M., GERIG, G., JOSHI, S., AND PIZER, S. 2003. Automatic and Robust Computation of 3D Medial Models Incorporating Object Variability. *Int. J. Comput. Vision* 55, 2, 107–122.

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*, 39–50.

SVENSSON, S. 2001. Reversible surface skeletons of 3d objects by iterative thinning of distance transforms. *Digital and image geometry: advanced lectures*, 400–411.

TAM, R., AND HEIDRICH, W. 2003. Shape Simplification Based on the Medial Axis Transform. In *IEEE Visualization 2003*.

WARD, A. D., AND HAMARNEH, G. 2009. The Groupwise Medial Axis Transform for Fuzzy Skeletonization and Pruning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 99.

Appendix

Guarantees for Conversion to Union of Balls. One of the free parameters of our method that need to be specified by the user is the tolerance $\delta > 0$ that controls the approximation of an input shape by a union of balls. Here we provide the following theoretical guarantees regarding this parameter:

1. For $\delta \rightarrow 0$, the Hausdorff distance between the smooth input shape and its approximation goes to zero. We can even provide bounds on the speed of convergence. These bounds depend on the shape, or more specifically on the so called *reach* of the shape.
2. For small values of δ the approximation becomes homotopy equivalent to the smooth input shape. Again, 'small' here is measured with respect to the reach of the input shape.
3. The number of balls in the approximation can be asymptotically bounded in terms of δ and the input shape.

To make our claims precise we need the following definitions. A *shape* can be any bounded, open set $S \subset \mathbb{R}^3$. The *boundary* of shape S is denoted as ∂S , the *local feature size* $\text{lfs}(x)$ maps any point $x \in \partial S$ to its distance from the medial axis of S , and finally, the *reach* is the minimum of the local feature size, i.e., $\text{rc}(S) = \min_{x \in \partial S} \text{lfs}(x)$ is a shape dependent constant.

Corollary 1. *For any open set $S \subset \mathbb{R}^3$ with smooth boundary and $\delta < \varepsilon \cdot \text{rc}(S)$ with $\varepsilon < 0.064$, the output U of our algorithm satisfies*

- i) U is homotopy equivalent to S
- ii) for any $x \in \partial U$, there exists a $y \in \partial S$ with $d(x, y) < \delta \cdot O(\varepsilon^2)$
- iii) the number of balls in U is limited by $O(\iint_S \frac{dx}{\delta^2})$

Proof Sketch. The proof is based on already proven properties of the sampling algorithm and polar balls. We use the algorithm of [Boissonnat and Oudot 2005] to generate a set P of sample points using the sizing field $\sigma(x) = \delta < \varepsilon \cdot \text{lfs}(S)$ for any $x \in \mathbb{R}^3$, see [Boissonnat and Oudot 2005] for details. The claim on the number of balls needed in the approximation (iii) follows from Lemma 8.3 of [Boissonnat and Oudot 2005] that provides an asymptotic upper bound of $O(\iint_S \frac{dx}{\delta^2})$ on the size of P . Since P limits the number of polar balls this implies (iii) Furthermore, from Lemma 8.2 of [Boissonnat and Oudot 2005] it follows that P is a so-called 'loose ε -sample', which is also an ε' -sample with $\varepsilon' = \varepsilon(1 + 8.5\varepsilon)$ according to Corollary 4.10 of [Boissonnat and Oudot 2005]. Being an ε' -sample means that any $x \in \partial S$ has a point in P at distance at most $\varepsilon \cdot \text{lfs}(x)$. Since δ is chosen such that ε' is less than 0.1, then conditions of Theorem 14 of [Amenta and Kolluri 2000] and Theorem 21 of [Amenta et al. 2001] are satisfied. These theorems immediately imply our claim of topological equivalence (i) and our claim of geometric closeness (ii) for the input shape and its approximation. \square

We have shown that U and S are topologically equivalent, therefore their medial axes have the same homotopy type, as well. But the geometric closeness on the boundaries of U and S does not imply closeness on their medial axes. However, a Hausdorff bound on the shape boundaries as in (ii) implies that certain stable subsets of the two medial axes, namely the λ -medial axes, are close to each other. This means that the medial axis of U captures all the features of the original shape S up to a size that is determined by δ . And the additional unstable parts of the medial axis of U will be removed in the scale-adaptive simplification step later in our pipeline.




Figure 12: A subset of the 380 discrete scale axes for the shape database of [Chen et al. 2009], all computed using the same settings $\delta = 0.01$ and $s = 1.1$. In the top left corner we show result for a shape described as point set. In the bottom right corner we show the discrete scale axis ($\delta = 0.005$, $s = 1.1$) for a scan of a *Cardium Pseudolima* shell model consisting of 1,268,908 triangles.