Pacific Graphics 2010
P. Alliez, K. Bala, and K. Zhou
(Guest Editors)

Volume 29 (2010), Number 7

A simple and robust thinning algorithm on cell complexes

L. Liu[1], E. W. Chambers[2], D. Letscher[2], and T. Ju[1]

[1]Washington University in St. Louis, USA
[2]St. Louis University, USA

Abstract

Thinning is a commonly used approach for computing skeleton descriptors. Traditional thinning algorithms often have a simple, iterative structure, yet producing skeletons that are overly sensitive to boundary perturbations. We present a novel thinning algorithm, operating on objects represented as cell complexes, that preserves the simplicity of typical thinning algorithms but generates skeletons that more robustly capture global shape features. Our key insight is formulating a skeleton significance measure, called medial persistence*, which identify skeleton geometry at various dimensions (e.g., curves or surfaces) that represent object parts with different anisotropic elongations (e.g., tubes or plates). The measure is generally defined in any dimensions, and can be easily computed using a single thinning pass. Guided by medial persistence, our algorithm produces a family of topology and shape preserving skeletons whose shape and composition can be flexible controlled by desired level of medial persistence.*

1. Introduction

Skeletons are common shape descriptors in vision and graphics, and are widely used in applications such as object segmentation, matching and retrieval, and animation. Thinning is a classical approach for computing skeletons. The concept originates from an illustrative definition of the medial axes, a construct closely related to skeletons: set the object boundary on fire and the medial axes is formed by the loci where the fire fronts meet and extinguish each other [Blu67]. Thinning simulates the fire-front propagation in a discrete manner, iteratively removing boundary elements of the object (like peeling off layers of an onion) until a thin piece is left. To produce skeletons useful for downstream applications as mentioned above, a thinning algorithm needs to retain important properties of the object such as its connectivity and major components.

Thinning algorithms are usually developed on objects represented as *digital pictures* [KR89], consisting of object points on a spatial grid. Traditional thinning algorithms consider local neighborhood of each point on the grid to identify points critical for preserving the topology or shape of the object. While simple to implement, these local criteria are very sensitive to small boundary perturbations (e.g., bumps), easily producing skeletons with spurious pieces. While more global measures can be used to identify meaningful skeleton parts during thinning, these measures are typically costly to compute (see next section for a brief review).

We introduce a new thinning algorithm that, on one hand, preserves the simplicity of a typical digital thinning algorithm, and on the other hand, yields stable and controllable skeletons that capture global shape features. The algorithm is developed on a discrete representation of objects as *cell complexes* [ZJH07]. Our key observation is formulating a global skeleton significance measure that can be easily computed by a single pass of iterative cell removal on the cell complex. This measure, called *medial persistence (MP)*, records the duration in which a discrete cell (e.g., an edge) persists in an isolated form (e.g., not adjacent to a face) during the removal process. We observed that faces and edges with high MP values lie respectively centered in significant plate-like and tube-like object parts (see Figure 1 (b,c) for an example). Guided by the measure, our algorithm can produce a family of surface and curve skeletons that robustly depict the object's topology and shape (see Figure 1 (d,e)).

Contributions Compared to existing thinning algorithms, our algorithm has the following unique set of features:

- **Simple:** The core of the algorithm is a simple iterative thinning procedure involving only local removal criteria. The same algorithm can be applied without changes to

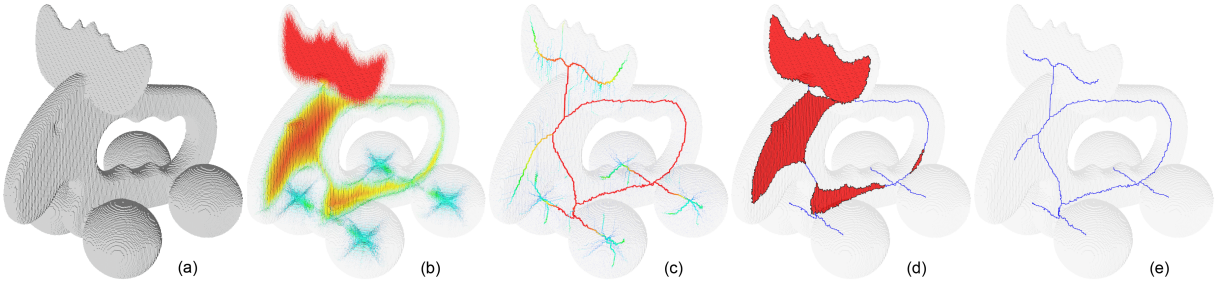

Figure 1: *To compute a skeleton of a 3D cell complex (a), our algorithm is guided by a skeleton significance measure,* medial persistence (MP)*, which highlights faces (b) and edges (c) of the complex that lie respectively centered to plate-like and tube-like object parts (these faces and edges are shown with redder color and in more contiguous pieces; see Section 4.2 for details on the rendering style). By varying the desired level of MP, the algorithm produces a family of skeletons containing surfaces and/or curves (d,e). Both MP and the skeletons are computed by simple, iterative cell removals.*

cell complexes in any dimensions, and to produce skeletons of any dimensions (e.g., curve or surface skeletons).

- **Robust:** The resulting skeletons faithfully capture anisotropic shape elongations (e.g., tubular and plate-like parts) by geometry at corresponding dimensions (e.g., skeletal curves and surfaces). The skeletons are stable under minor boundary perturbation or changes in the resolution of the cell complex.
- **Controllable:** The shape and dimensionality of the skeleton can be flexibly controlled by a small number of parameters. Varying the parameters produces a family of topology-preserving skeletons with different composition of surfaces and curves.

2. Previous work

There is a sizable literature on computing skeletons, and we refer readers to a recent survey book [SP08] and articles [SB98, CM07] for extensive reviews. For the purpose of this work, we will focus on methods related to thinning.

2.1. Thinning on digital pictures

Thinning on 2D and 3D objects represented as disjoint lattice points (i.e., *digital pictures*) has been extensively studied in the area of *digital topology* [Ros79]. A typical thinning scheme involves, at each iteration, identifying boundary points and deleting those whose removal would not introduce a topological change (e.g., breaking or piercing) or loss of geometric features (e.g., shrinking of curves or surfaces). The discrete setting makes it is possible to identifying these points based on a point's local neighborhood on the lattice [Ber94, Ber95, PK99]. Employing local criteria makes the thinning algorithm straight-forward to implement, and numerous algorithms were developed based on this simple scheme (see survey and reviews in [LLS92, Pal08, SP08]). While simple to implement, these local criteria are very sensitive to small perturbations on the object's boundary. As a result, thinning based on local criteria often leads to skeletons containing spurious branches that do not reflect global shape features (see examples in Figure 8 on the last page), and post-process pruning [SB98, JBC07] is required.

To improve the stability of skeletons, thinning can be guided by criteria that consider more global properties of the shape. These criteria, which we call *skeleton significance measures*, can be obtained by functions over the entire object such as distance fields [BNSdB99], vector fields [SBTZ02], and force fields [CTK00, BB08], or by geometric relations between the skeleton point and the object boundary, such as the distance or angle formed by a skeleton point to its closest boundary points [Blu67, ACK01, DZ02, DDS03, SFM05], and the surface geodesic metrics between these boundary points [OK95, DS06, RvWT08].

Although these global measures can significantly improve the stability of skeletons, their computation is typically costly and involves global operations at each point being considered. We also point out that most of these measures are designed for skeletons at a specific dimension, such as a surface skeleton or a curve skeleton in 3D. In contrast, our significance measure (medial persistence) has a unified formulation for skeleton elements at different dimensions (e.g., edges or faces), and the measured values at all skeleton elements can be easily computed by a single thinning pass.

2.2. Thinning on other model representations

The disjoint nature of points in a digital picture can create obstacles for topology and geometry analysis [Ros79]. Such difficulty has led several researchers to explore alternative discrete representations as media for thinning. Kovalevsky [Kov04], Wiederhold and Morales [WM08] consider an *abstract cell complex*, a connected and possibly open (boundary-less) set consisting of cells at various dimensions. Theoretical properties of the representation are presented and a topology-preserving thinning algorithm is proposed. However, the algorithm still results in noisy skeletons as it uses local criteria, similar to those in thinning algorithms on digital pictures, for deciding what cells should be preserved during thinning.

In this work, we consider a more restricted type of cell complex that is a *closed* set of cells, which was first considered by Zhou and colleagues [ZJH07, JZH07] in the context of topology repair of 3D models. In these works, the original complex is reduced, using a topology-preserving cell removal operator (to be detailed in Section 3), to a minimal skeleton that only preserves the object's topology (e.g., a banana will be reduced to a point). The skeleton is subsequently used for modifying the topological structure of the object. With a different aim in mind, we consider in this work the same cell complex representation and the cell removal operator as in [ZJH07, JZH07], but focus on developing a robust shape-preserving criteria to protect skeletal elements during thinning that are important to capture the object's shape, in addition to its topology.

3. Cell complexes and cell removal

We start by briefly reviewing the cell complex representation and an associated removal operator, which are first considered by Zhou *et. al.* [ZJH07] for topology-preserving thinning. Unlike digital pictures, a cell complex consists of a connected set of explicit geometric elements (e.g., edges and faces) that are ideal for representing skeletons at different dimensions (e.g., curve and surface skeletons).

3.1. Cell complexes

A cell complex is a closed set of k-cells, each homeomorphic to an open ball in k-dimensions. For example, a point is a 0-cell, an edge without its end points is a 1-cell, and a polygon without its border is a 2-cell. By definition, if a cell δ is in a cell complex S, all cells on the border of δ, called the *facet* of δ, are also in S. A 2D example of a cell complex is shown in Figure 2 (a), where a cell δ (a triangle) and one of its facets σ (an edge) are marked.

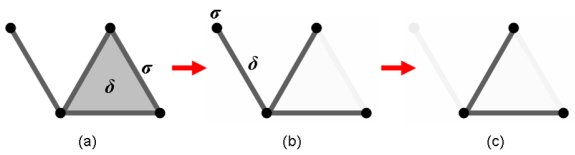

Figure 2: *A cell complex (a), and results after one (b) and two (c) simple pair removals.*

A cell complex can be constructed in several ways. If the input object is already represented as a digital picture (which can be converted from other representations using voxelization techniques such as [Ju04]), we can build the complex by creating a 0-cell for each object point on the grid, a 1-cell for each grid edge connecting two object points, a 2-cell for each grid face whose corners are all object points, and so on. Alternatively, cell complexes can be built from any spatial decomposition, such as the triangulation of a 2D polygon or the tetrahedralization of a 3D polyhedron.

3.2. A removal operator

To perform thinning on cell complexes, Zhou *et. al.* [ZJH07] consider a pair-wise cell removal operator, which is a counterpart to point-removal in digital thinning. Let S be a cell complex, we define:

Definition 1 A *simple pair* is a pair of cells $\{\delta, \sigma\} \in S$ such that σ is a facet of δ, and that δ is the only cell of which σ is a facet. In a simple pair, δ is called the *simple cell*, and σ is the *witness facet*.

It was shown in [ZJH07] that removing a simple pair from a cell complex S results in a sub-complex S' that is a valid cell complex and homotopy equivalent to S. An example of removing simple pairs in shown in Figure 2. Note that the complexes in (b,c) preserve the topology of that in (a).

4. Medial persistence

While removing simple pairs preserves the topology of a cell complex, exhaustive removal may cause significant loss of shape features. For shape description, thinning needs to additionally identify and preserve those cells that carry important shape information. Note that different kinds of cells are good at describing different shape features. For example, 1-cells (edges) on a skeleton curve depict well tubular shapes, while 2-cells (faces) on a skeleton surface mimic plate-like shapes. Intuitively, a k-cell on the skeleton describes well a shape that has *anisotropic* elongation in k directions (e.g., $k = 1$ for tubes and $k = 2$ for plates).

In the following, we will develop a significance measure to guide thinning, which computes for each k-cell on the skeleton the amount of k-directional anisotropy of the surrounding shape. The measure itself is motivated and defined by an iterative thinning process, which we will explain first.

4.1. Iterative thinning

Consider the following iterative cell removal process. At each iteration, we identify *all* simple cells (according to Definition 1) at *all* dimensions, and remove them simultaneously from the complex together with their witness facets. If a simple cell has multiple facets that are candidates as witness facets, an arbitrary one is selected.

This process is illustrated on a 2D square cell complex in Figure 3 top. Observe that each iteration "strips off" a layer of the object, due to the simultaneous removal of simple cells that lie at the boundary of the object. As the iterations proceed, the 2D object is shrunk to some 1D skeleton curves, which continue to be eroded from their ends. Notice also that the object's topology is retained under the simultaneous removal of simple pairs. To see why, it suffices to note that the simple pairs selected to be removed at each iteration are disjoint from each other, and that each pair stays "simple" when other simple pairs are removed.

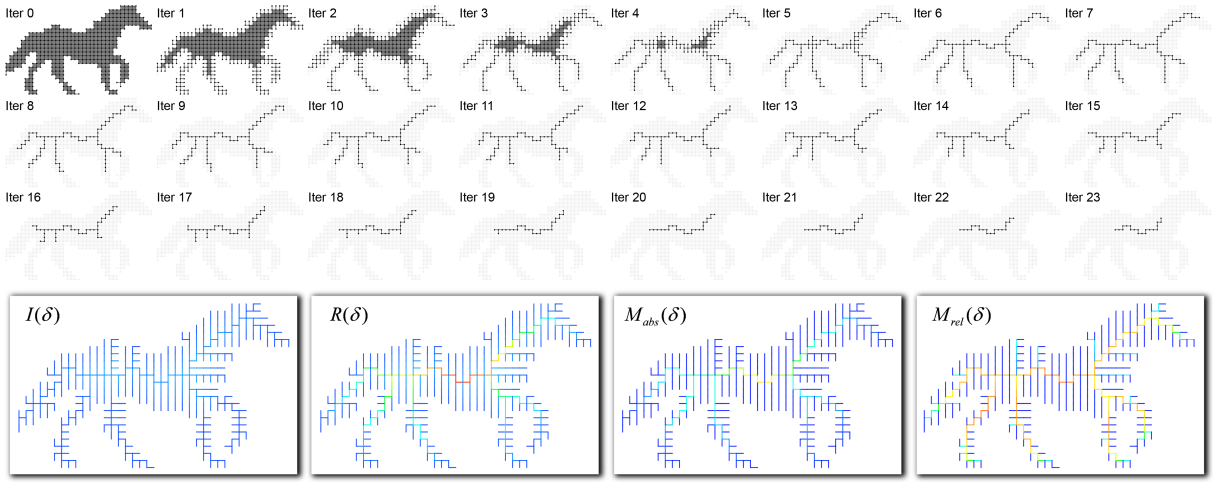

Figure 3: *Top: first 23 iterations during the iterative thinning on a 2D cell complex. Bottom: isolation iteration ($I(\delta)$), removal iteration ($R(\delta)$), and medial persistence ($M_{abs}(\delta), M_{rel}(\delta)$) for all isolated edges δ during thinning (visualized on a heat color scale) .*

4.2. Formulating medial persistence

Motivation Here is the key observation that motivates our formulation: a k-cell that lies centered to a shape part with prominent k-directional anisotropy tends to *persist* in an isolated form during the above iterative thinning process. Here, we say a k-cell is *isolated* if it is not a facet to any $(k+1)$-cell. Consider again the 2D thinning process shown in Figure 3 top. Note that the edges centered at the elongated parts of the object (e.g., legs, tail, neck and torso) remain isolated in the complex for many more iterations during thinning than other edges in the cell complex. Likewise, when thinning on a 3D cell complex, we observed that faces centered at flat and wide parts, and edges centered at long and thin parts, remain isolated for more iterations than other faces and edges (see the accompanying video).

The observation can be explained intuitively as follows. Let $I(\delta)$ be the earliest thinning iteration after which a k-cell δ is isolated, and $R(\delta)$ be the iteration at which δ is removed. In an approximate sense, $I(\delta)$ and $R(\delta)$ respectively measure the maximum *isotropic* elongation of shape in $k+1$ and k directions around δ. For example, at an edge δ of a 2D complex, $I(\delta)$ measures the shortest (discrete) distance from δ to the object boundary, which depicts the size of the maximum disk centered at δ and inscribed in the object. On the other hand, $R(\delta)$ measures the longest distance from δ to the object boundary *along* the skeleton curves, which depicts how far the object stretches sideways from δ. As a result, the duration in which δ remains isolated in thinning, which is the difference between $I(\delta)$ and $R(\delta)$, reflects how much *anisotropic* elongation the shape has in k directions.

Formulation Based on the observation, we formulate a skeleton significance measure, called *medial persistence (MP)*, as the duration in which a cell remains isolated during thinning. Specifically, for a k-cell δ that is isolated at some iteration during iterative thinning, we compute two scores to capture the absolute and relative duration of isolation:

$$\begin{aligned} M_{abs}(\delta) &= R(\delta) - I(\delta) \\ M_{rel}(\delta) &= 1 - \frac{I(\delta)}{R(\delta)} \end{aligned}$$

Properties We note that $I(\delta) \leqslant R(\delta)$, hence both MP scores are non-negative and $M_{rel}(\delta) \leqslant 1$. Some cells may never be removed by thinning if it is critical for retaining the topology of the object. For example, iterative thinning on a torus will terminate with a closed curve, while thinning on a ball with an internal cavity will result in a closed surface. For a cell δ that remains after thinning, we intuitively set $R(\delta) = \infty$, in which case both MP scores would reach their maximum.

The MP scores are defined and computed in a dimension-oblivious manner. They are suited to rate skeleton geometry at any k dimensions for a d-dimensional complex where $k < d$. Regardless of the dimension of the complex, both measures M_{abs}, M_{rel} for all cells at dimension $k < d$ can be computed by a single pass of iterative thinning.

Examples As a 2D example, Figure 3 bottom plots the isolation iteration, removal iteration, and MP scores for all isolated edges during the thinning process at the top. Observe that while M_{abs} scores higher in parts with greater absolute length (e.g., torso versus legs), M_{rel} describes elongation in a scale-invariant way. The scale-invariance of M_{rel} also implies that it may evaluate high on small perturbations on the boundary, where M_{abs} would be low. Therefore cells with high values in both scores would represent parts that are both sharp in shape and large in size.

Figure 4 (b) shows the MP scores in 3D for faces and edges resulted from thinning the cell complex in (a). To facilitate visualization, the faces and edges (in this and other

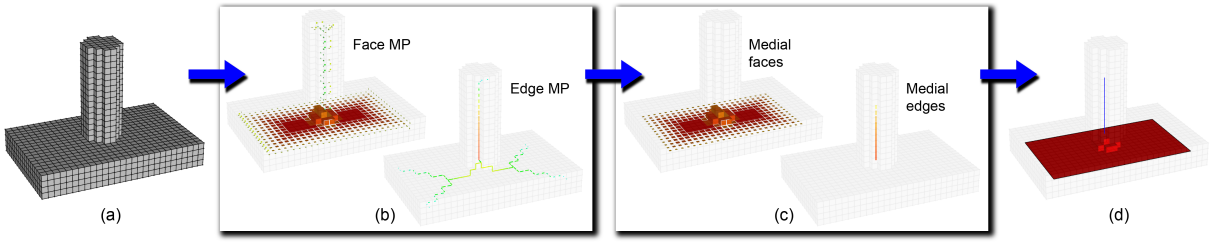

Figure 4: *Algorithm flow: an input cell complex (a), medial persistence computed after one pass of thinning (faces and edges colored by their M_{rel} scores and scaled by their M_{abs} scores) (b), clustered medial edges and faces (c), and the final skeleton (curves and surfaces shown in blue and red) after another pass of thinning that preserves the medial edges and faces (d).*

figures) are colored by their M_{rel} scores and drawn at a scale proportional to their M_{abs} scores. Observe that faces and edges with high scores in both M_{rel} and M_{abs} (shown in (c)) lie respectively at visually prominent plate-like and tube-like parts of the object.

Comparisons A distinctive feature of medial persistence is that it captures the "aspect ratio" of shape features rather than their sizes. This is illustrated well in Figure 4 (b) right, where the edge MP score M_{rel} (indicated by coloring) is much higher on the centerline of the cylinder, which exhibits a dominant elongation in 1 dimension, than on the centerlines of the bottom plate, which has a rather uniform elongation in 2 dimensions, even though the plate appears to be larger than the cylinder. In contrast, a standard pruning approach based on lengths of skeleton branches would favor the centerlines of the plate, which appear longer than the centerline of the cylinder.

There has been few global skeleton significance measures proposed for 3D skeletons. Here we compare our medial persistence with a recent measure of Reniers et al. [RvWT08] in Figure 5. Reniers extended the Feature Distance measure in 2D [OK95], which expresses the length of the shortest curve on the shape boundary between the closest boundary points to a skeleton point, to evaluate 3D surface skeletons using lengths of geodesic curves and 3D curve skeletons using approximated areas of geodesic patches. Observe from Figure 5 (d,e) that FD measures favor regions on the skeleton closer to the center of the object. In comparison, the medial persistence highlights regions with high anisotropy, such as flat and wide regions (e.g., wings and tail) and long and thin parts (e.g., wings and tail trunk), even though they may be located away from the object center.

5. The algorithm

Guided by the medial persistence measure, it is straightforward to compute a topology-preserving and shape-depicting skeleton. All it needs is to perform the same iterative thinning process (Section 4.1) while retaining cells with high MP scores. It is tempting to combine the two tasks, computing MP and retaining cells, within a single thinning pass. However, retaining cells during thinning may have an impact on the MP values of other cells computed at later iterations. As a result, we split the tasks into two separate thinning passes, and insert an optional clustering stage to improve the skeleton contiguity. The algorithm proceeds as follows for a d-dimensional cell complex (see Figure 4):

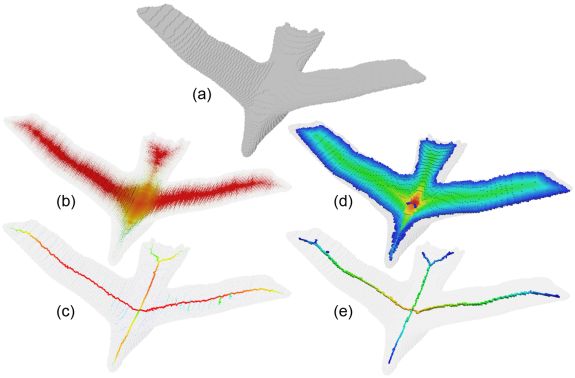

Figure 5: *Compare medial persistence on faces (b) and edges (c) in a cell complex with the extended Feature Distance measure [Reniers et al. 2008] on skeleton surfaces (d) and curves (e).*

Step 1 (Thinning): Perform iterative thinning on the input cell complex and compute scores M_{abs}, M_{rel} for all cells (Figure 4 (b)).

Step 2 (Clustering): For each $k < d$, identify connected components, whose sizes are greater than some threshold τ^k, of k-cells that score higher than thresholds $\epsilon^k_{abs}, \epsilon^k_{rel}$. The clustered cells are called *medial cells* (Figure 4 (c)).

Step 3 (Thinning): Perform iterative thinning again on the original complex while maintaining the medial cells (Figure 4 (d)).

Note that the algorithm results in a skeleton that is guaranteed to be homotopy equivalent of the input complex. Even though the collection of medial cells found by the second step can exhibit a different topology than the object, the final thinning step enforces the topology-preservation. The final skeleton is composed of the medial cells and a minimal set of cells necessary to maintain the connectivity and topological structure of the input complex.

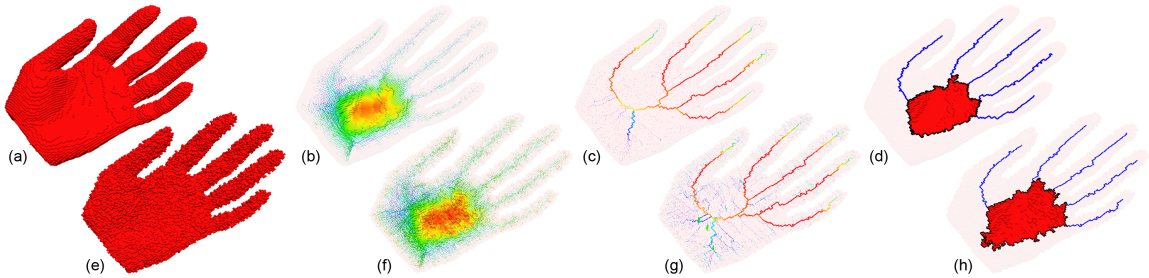

Figure 6: *The result of our algorithm on an original model (a) and one with synthetically added noise (e), showing MP scores on faces (b,f) and edges (c,g) (in the same rendering style as Figure 4 (b)) and the surface-and-curve skeletons (d,h).*

The composition of the skeleton can be intuitively controlled using the provided parameters. In particular, by setting the absolute MP threshold $\varepsilon^k_{abs} = \infty$ for some $k < d$, the algorithm would produce a skeleton that is void of k-dimensional geometry, unless such geometry is critical for preserving the topology. For example, setting $\varepsilon^2_{abs} = \infty$ when thinning a 3D cell complex would result in a curve-only skeleton except for complexes with internal cavities, which would additionally have closed surfaces in their skeletons.

6. Results

We demonstrate our algorithm on computing skeletons for a suite of 3D models. The input cell complexes are converted from triangular meshes as discussed in Section 3.1. We show skeletons generated at two parameter settings: surface-and-curve skeletons at $\varepsilon^k_{abs} = 0.05L, \varepsilon^k_{rel} = 0.5$ for both $k = 1,2$ where L is the width of the bounding box, and curve-only skeletons at the same setting except $\varepsilon^2_{abs} = \infty$. The minimum component size is $\tau^k = (0.05L)^k$ for $k = 1,2$ in all examples.

We first examine the stability of the MP scores and skeletons under boundary perturbation and varying discretization levels. In Figure 6, we compare the result on a hand model (a) and a synthetically damaged model (e) produced by applying two iterations of thinning on (a) during which simple pairs are randomly removed. Observe that, despite the severe boundary perturbation, the distribution of MP scores is not significantly affected. Faces and edges with high MP scores in both models faithfully depict the plate-like parts (e.g., the palm) and tubular parts (e.g., the fingers) of the object, yielding skeletons with similar structures. Figure 7 tests our algorithm on a same model voxelized at increasing resolutions (using the Fish model shown in Figure 8 top). Notice that the resulting skeletons, computed with the same parameter setting, maintain a stable structure.

We present a gallery of examples in Figure 8. Observe that faces (2nd row) and edges (3rd row) with high MP scores capture well the plate-like and tube-like parts of these models. We compare our skeletons with curve skeletons generated by a typical thinning algorithm on digital pictures [PK99] (bottom row). Both our algorithm and [PK99] involve a simple iterative thinning process with local removal

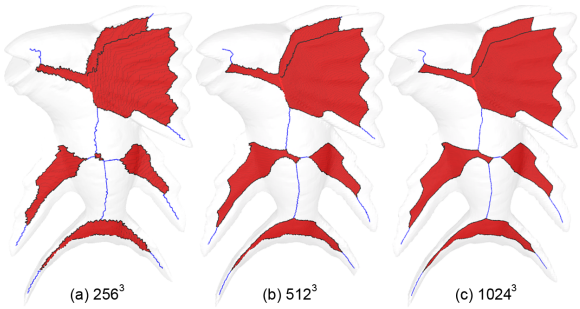

Figure 7: *Skeletons computed on cell complexes at increasing grid resolutions.*

criteria. Notice that the results of the latter often contain spurious branches (which are typical for digital thinning algorithms of comparable complexity), while ours are more stable due to the guidance by medial persistence. The running times and memory footprints of our algorithm for these models are shown at the bottom of the 5th row.

7. Conclusion and discussions

In this paper, we present a simple and robust algorithm to compute discrete skeleton descriptors by iterative thinning. The algorithm shares a similar structure as most thinning algorithm on digital pictures, but generates skeletons that more stably capture global shape features. This is achieved by formulating a novel skeleton significance measure (medial persistence) that captures anisotropic shape elongations and is simple to compute. The algorithm additionally offers flexible control over the shape and composition of the resulting skeletons.

Discussions There are a number of aspects of our algorithm that leave room for further improvement. First, we use an arbitrary, fixed ordering when selecting from multiple witness facets of a simple cell for removal. Although we have not noticed any significant difference in results with different orderings, further analysis on the effect of ordering is in order. Second, as our thinning follows the grid structure of the cell complex, the resulting skeletons may be biased by the anisotropy of the grid in addition to that of the ob-

ject. We have noticed in our tests that applying our algorithm to irregularly arranged cells (e.g., an octree) does not result in skeletons that represent the shape in a meaningful way. Even on cells arranged as regularly as a cubic complex, the MP scores may still be affected by the grid orientation (e.g., along the diagonals of the wheels in Figure 1 (b,c)). To this end, we are currently exploring ideas such as distance-guided removal of simple pairs at each iteration. Third, due to the iterative and parallel nature of simple pair removals, our algorithm can be significantly accelerated by a parallel, hardware-accelerated implementation.

Another promising direction of future research is investigating a continuous formulation of our discrete thinning procedure and the medial persistence measure. In particular, the iterative thinning described in Section 4.1 mimics a continuous erosion of a manifold simultaneously from all open boundaries (corresponding to simultaneous simple-pair removals), during which lower-dimensional manifolds are formed (corresponding to a cell becoming isolated) when the erosion fronts along a higher-dimensional manifold meet. In this analogy, medial persistence becomes the difference in the arrival time of erosion fronts on manifolds at different dimensions. We believe such a continuous formulation exists, as suggested by the results when running the algorithm at increasing discretization levels (see Figure 7).

Acknowledgement The 3D models in this paper are downloaded from Aim@Shape online model repository. This work is supported in part by NSF grants IIS-0705538, CCF-0702662, and IIS-0846072.

References

[ACK01] AMENTA N., CHOI S., KOLLURI R. K.: The power crust. In *SMA '01: Proceedings of the sixth ACM symposium on Solid modeling and applications* (2001), pp. 249–266.

[BB08] BRUNNER D., BRUNNETT G.: Fast force field approximation and its application to skeletonization of discrete 3D objects. *Computer Graphics forum journal 27*, 2 (2008), 261–270.

[Ber94] BERTRAND G.: Simple points, topological numbers and geodesic neighborhoods in cubic grids. *Pattern Recogn. Lett. 15*, 10 (1994), 1003–1011.

[Ber95] BERTRAND G.: A parallel thinning algorithm for medial surfaces. *Pattern Recogn. Lett. 16*, 9 (1995), 979–986.

[Blu67] BLUM H.: A transformation for extracting new descriptors of form. *Models for the Perception of Speech and Visual Form* (1967), 362–80.

[BNSdB99] BORGEFORS G., NYSTRÖM I., SANNITI DI BAJA G.: Computing skeletons in three dimensions. *Pattern Recognition 32*, 7 (1999), 1225–1236.

[CM07] CORNEA N. D., MIN P.: Curve-skeleton properties, applications, and algorithms. *IEEE Transactions on Visualization and Computer Graphics 13*, 3 (2007), 530–548. Member-Silver,, Deborah.

[CTK00] CHUANG J.-H., TSAI C.-H., KO M.-C.: Skeletonization of three-dimensional object using generalized potential field. *IEEE Trans. Pattern Anal. Mach. Intell. 22*, 11 (2000), 1241–1251.

[DDS03] DIMITROV P., DAMON J. N., SIDDIQI K.: Flux invariants for shape. In *In International Conference on Computer Vision and Pattern Recognition* (2003), pp. 835–841.

[DS06] DEY T. K., SUN J.: Defining and computing curve-skeletons with medial geodesic function. In *SGP '06: Proceedings of the fourth Eurographics symposium on Geometry processing* (Aire-la-Ville, Switzerland, Switzerland, 2006), Eurographics Association, pp. 143–152.

[DZ02] DEY T. K., ZHAO W.: Approximate medial axis as a voronoi subcomplex. In *SMA '02: Proceedings of the seventh ACM symposium on Solid modeling and applications* (2002), pp. 356–366.

[JBC07] JU T., BAKER M. L., CHIU W.: Computing a family of skeletons of volumetric models for shape description. *Comput. Aided Des. 39*, 5 (2007), 352–360.

[Ju04] JU T.: Robust repair of polygonal models. *ACM Trans. Graph. 23*, 3 (2004), 888–895.

[JZH07] JU T., ZHOU Q.-Y., HU S.-M.: Editing the topology of 3d models by sketching. *ACM Trans. Graph. 26*, 3 (2007), 42.

[Kov04] KOVALEVSKY V.: Algorithms in digital geometry based on cellular topology. In *IWCIA* (2004), pp. 366–393.

[KR89] KONG T. Y., ROSENFELD A.: Digital topology: introduction and survey. *Comput. Vision Graph. Image Process. 48*, 3 (1989), 357–393.

[LLS92] LAM L., LEE S.-W., SUEN C. Y.: Thinning methodologies-a comprehensive survey. *IEEE Trans. Pattern Anal. Mach. Intell. 14*, 9 (1992), 869–885.

[OK95] OGNIEWICZ R. L., KÜBLER O.: Hierarchic voronoi skeletons. *Pattern Recognition 28*, 3 (1995), 343–359.

[Pal08] PALAGYI K.: A 3d fully parallel surface-thinning algorithm. *Theor. Comput. Sci. 406*, 1-2 (2008), 119–135.

[PK99] PALAGYI K., KUBA A.: A parallel 3d 12-subiteration thinning algorithm. *Graph. Models Image Process. 61*, 4 (1999), 199–221.

[Ros79] ROSENFELD A.: Digital topology. *The American Mathematical Monthly 86*, 8 (1979), 621–630.

[RvWT08] RENIERS D., VAN WIJK J., TELEA A.: Computing multiscale curve and surface skeletons of genus 0 shapes using a global importance measure. *IEEE Transactions on Visualization and Computer Graphics 14*, 2 (2008), 355–368.

[SB98] SHAKED D., BRUCKSTEIN A. M.: Pruning medial axes. *Comput. Vis. Image Underst. 69*, 2 (1998), 156–169.

[SBTZ02] SIDDIQI K., BOUIX S., TANNENBAUM A., ZUCKER S. W.: Hamilton-jacobi skeletons. *International Journal of Computer Vision 48*, 3 (July 2002), 215–231.

[SFM05] SUD A., FOSKEY M., MANOCHA D.: Homotopy-preserving medial axis simplification. In *SPM '05: Proceedings of the 2005 ACM symposium on Solid and physical modeling* (New York, NY, USA, 2005), ACM, pp. 39–50.

[SP08] SIDDIQI K., PIZER S. M.: *Medial Representations*. Springer, 2008.

[WM08] WIEDERHOLD P., MORALES S.: Thinning on quadratic, triangular, and hexagonal cell complexes. In *IWCIA* (2008), pp. 13–25.

[ZJH07] ZHOU Q.-Y., JU T., HU S.-M.: Topology repair of solid models using skeletons. *IEEE Transactions on Visualization and Computer Graphics 13*, 4 (2007), 675–685.

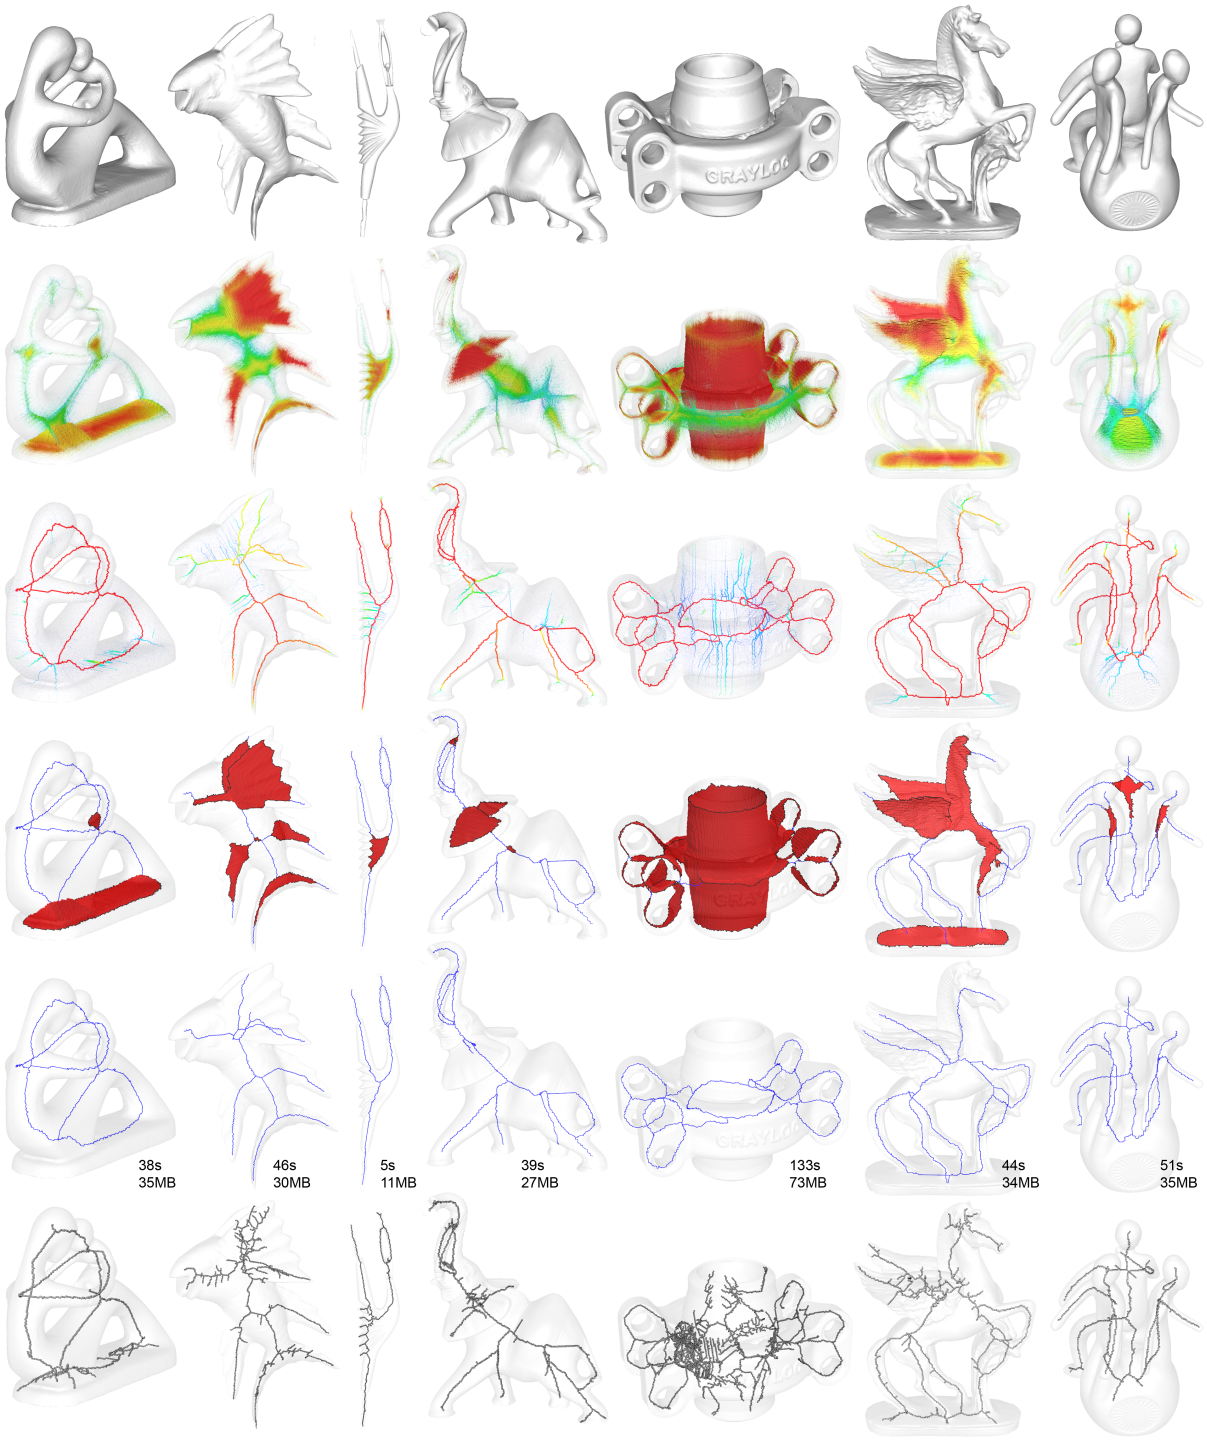

Figure 8: *Results on a gallery of 3D models (top row, all converted at resolution* 256^3*), showing MP scores for faces (2nd row) and edges (3rd row) (rendered in the same style as Figure 4 (b)), surface-and-curve skeletons (4th row) and curve-only skeletons (5th row), and curve skeletons computed by a digital thinning algorithm [Palagyi and Kuba 1999]. Time and memory used for computing curve-only skeletons are provided (on a PC with 2GB of main memory and 2.2GHz CPU).*