

Regularizing 3D Medial Axis Using Medial Scaffold Transforms

Ming-Ching Chang

mcchang@lems.brown.edu

LEMS, Engineering, Brown University, USA

Benjamin B. Kimia

kimia@lems.brown.edu

LEMS, Engineering, Brown University, USA

Abstract

This paper addresses a key bottleneck in the use of the 3D medial axis ($\mathcal{M}\mathcal{A}$) representation, namely, how the complex $\mathcal{M}\mathcal{A}$ structure can be regularized so that similar, within-category 3D shapes yield similar 3D $\mathcal{M}\mathcal{A}$ that are distinct from the non-category shapes. We rely on previous work which (i) constructs a hierarchical $\mathcal{M}\mathcal{A}$ hypergraph, the medial scaffold ($\mathcal{M}\mathcal{S}$), and (ii) the theoretical classification of the instabilities of this structure, or transitions (sudden topological changes due to a small perturbation). The shapes at transition point are degenerate. Our approach is to recognize the transitions which are close-by to a given shape and transform the shape to this transition point, and repeat until no close-by transitions exists. This move towards degeneracy is the basis of simplification of shape. We derive 11 transforms from 7 transitions and follow a greedy scheme in applying the transform. The results show that the simplified $\mathcal{M}\mathcal{A}$ preserves within-category similarity, thus indicating its potential use in various applications including shape analysis, manipulation, and matching.

1. Introduction

The *medial axis* ($\mathcal{M}\mathcal{A}$) of a shape is the closure of centers of maximal spheres that are at least tangent to the surface at two places. It is a *generic* representation useful in many 2D/3D shape modeling applications, including feature detection (ridges [17], flat/tubular regions [16]), animation [37], and shape simplification [33]. In computer vision, the $\mathcal{M}\mathcal{A}$ is used in shape matching [27, 5], retrieval [29, 18], morphological analysis, and segmentation [25, 9].

The $\mathcal{M}\mathcal{A}$ has great promise as a universal model for shape [20] since: (i) it is intuitively appealing in representing *elongated* and *branching* objects; (ii) a radius function along the $\mathcal{M}\mathcal{A}$ trace encodes the varying width of the shape; (iii) important global *features* such as curvature extrema (ridges), necks, limb, thin/thick parts are made explicit; (iv) a *hierarchy of scales* is built-in the representation; (v) it is *complete*, i.e., a full reconstruction is always possible [11]; and (vi) it provides a powerful framework to model and generate shapes [35, 22].

The $\mathcal{M}\mathcal{A}$ structure in 3D is more complex than in 2D

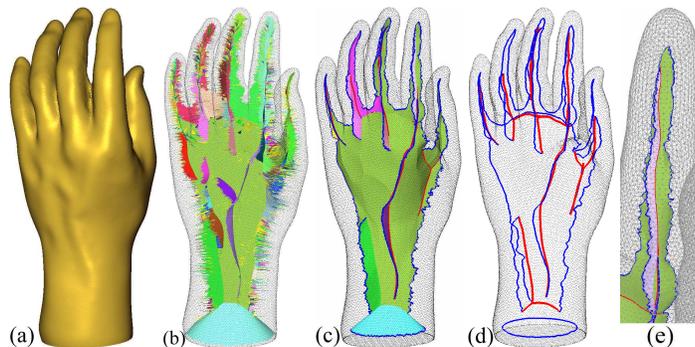


Figure 1. Regularizing the $\mathcal{M}\mathcal{A}$ of a scan of a hand ($\approx 38k$ points, from Polhemus) toward simplification. (a) The computed $\mathcal{M}\mathcal{S}$ and an associated meshing of the surface ($\approx 76k$ faces) and the initial noisy $\mathcal{M}\mathcal{S}$ (9,574 sheets) [6]. The $\mathcal{M}\mathcal{A}$ in (b) is greatly simplified via a series of transforms into a hypergraph of only 10 sheets, 33 curves, and 23 nodes in (c), reflecting the structure of the shape. (d) is the graph of (c) with sheets implicit, where the rib curves are shown in blue and axial curves are in red. The fore finger is zoomed on (e).

since it consists of connected sheets, curves, and points. Giblin and Kimia classify these points into one type of $\mathcal{M}\mathcal{A}$ sheet, two types of $\mathcal{M}\mathcal{A}$ curves, and two types of $\mathcal{M}\mathcal{A}$ points [12]. This classification has been used by Leymarie and Kimia [20] to propose the notion of a *medial scaffold* ($\mathcal{M}\mathcal{S}$), a hierarchical representation of the $\mathcal{M}\mathcal{A}$ points. Algorithms for computing the $\mathcal{M}\mathcal{S}$ from unorganized points have been developed [20, 6], which also effectively meshes the surface as well, Figure 1a,b. The $\mathcal{M}\mathcal{A}$ structure is complete, but is also complex and redundant with respect to a qualitative representation of the structure: as the shape is perturbed it undergoes *transitions*, or sudden changes in the $\mathcal{M}\mathcal{S}$ topology, so that highly similar shapes can have drastically different $\mathcal{M}\mathcal{S}$ topology. Giblin and Kimia classify these 3D transitions of the $\mathcal{M}\mathcal{S}$ into 7 types [14, 13], as shown in detail in Figure 2.

This paper addresses a key bottleneck in the use of the 3D $\mathcal{M}\mathcal{A}$ representation, namely, how the complex $\mathcal{M}\mathcal{A}$ structure can be regularized so that similar, within-category 3D shapes yield similar 3D $\mathcal{M}\mathcal{A}$ that are distinct from the non-category shapes. Our approach can be illustrated by relying on a simpler 2D example of Figure 3 [34, 27]: de-

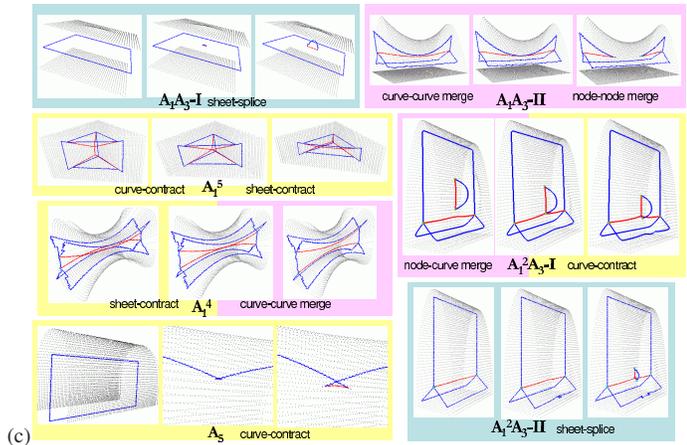
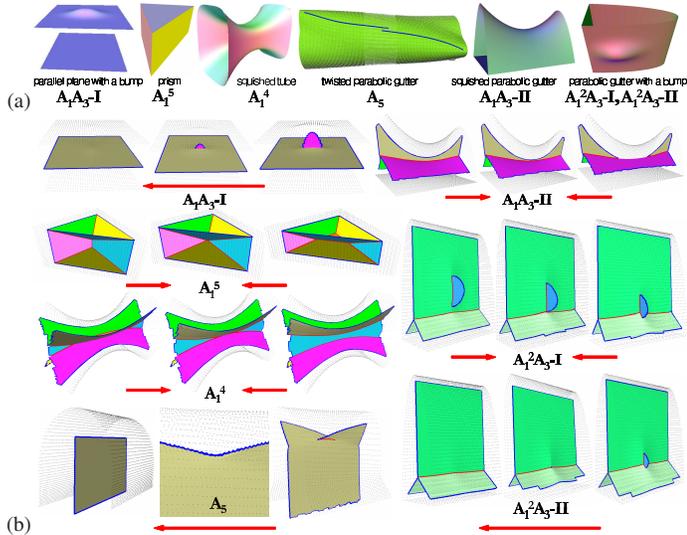


Figure 2. (a) Seven example shapes near the transition point across a one-parameter family of shape deformation [13]. (b) Simulation of the 7 cases of MS transitions. In (c) all medial sheets are hidden to better visualize their structures. The 11 MS transforms related to the 7 transitions are labelled as well.

generate shapes are simpler, so we transform all shapes to close-by more degenerate shapes as indicated by the near-transition MA configurations. The 7 transitions result in 11 transforms (red arrows in Figure 2b). These transforms are based on topological changes to the MA and to the shape, and described in detailed here. In fact, they are already based in computing the MS itself [21, 6].

Our contribution is as follows. First, we present a **dual-scale** representation of the MA topology and geometry, Figure 4, to allow for an explicit computation of the MA transition (Section 3). Second, we have identified 11 MS transforms covering all 7 transitions in Figure 2. In particular, traditional *pruning* of the MA can be recognized as the *splice* transform, but where our transform in addition to removing a medial sheet branch, also merges the remaining sheets. All other transforms (*contract*, *merge*) operate in the same spirit to fully simulate all transitions. Our

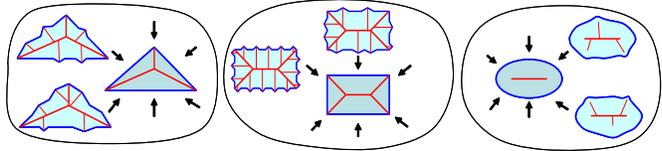


Figure 3. Stable MA regularization by grouping similar shapes under perturbation into equivalence classes and simplifying them toward a *representative* shape (center of arrows) of each class.

Coarse-scale: hypergraph

- Vertex: A_1^4 or A_1A_3 node
- Link: A_1^3 or A_3 curve
- Hyperlink: A_1^2 sheet

Fine-scale: (non-manifold) mesh

- Vertex: A_1^4 node element
- Edge: A_1^3 curve element
- Face: A_1^2 sheet element

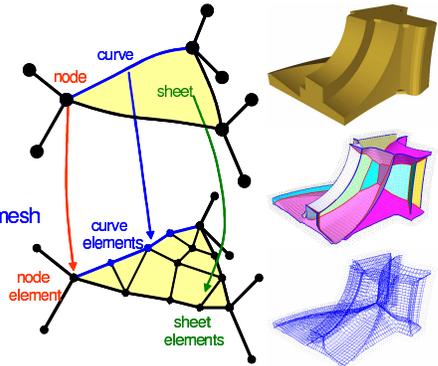


Figure 4. The proposed *dual-scale* MS representation to simulate the *true* MA of the underlying shape. The coarse-scale is a *hypergraph* capturing the *topology* and the fine-scale is a (not necessary manifold) *mesh* capturing the *geometry* and other information.

MS regularization process consists of a sequence of transforms rank-ordered by a saliency measure (Section 4) and performed in a *greedy* least-cost first fashion. Third, a **computational scheme** to realize the proposed framework in practice is also developed (Section 5). Our system takes input up to 300k points and computes the regularized MS efficiently and virtually with no manual parameters involved. We associate all sample points on the shape with their corresponding MS elements and maintain them consistently during all transforms. The advantage is a coupled *shape-skeleton* structure suitable for subsequent shape modeling applications (Section 6) ¹.

Our transform-based MA regularization is drastically different from the traditional approaches which focus on *pruning* medial sheets, which is tantamount to effecting two types of transforms (splicing out protrusions A_1A_3 -I and $A_1^2A_3$ -II) ². However, they are silent on structural inter-connectivity between the sheets which are illustrated in Figure 2 and do occur in practice all the time (see Section 2 for details). Our approach not only advocates a holistic component-based transformation in the first two cases (as opposed to element-based pruning), but also introduces **nine** additional cases which are required in an effective regularization of the MA so that the **qualitative** structure of

¹ The related work of *pair-mesh* [Shamir & Shaham, GM'06] also provides a coupled structure. However, a solid mesh is assumed (for the inner \mathcal{VD} and the PowerCrust) and it works in a different manner.

² The notation A_k^n indicates k -fold tangency at n points. A_1 is regular tangency.

the \mathcal{MA} emerges; see Figure 12 as an example.

Another advantage of computing a qualitative \mathcal{MA} structure is the computation of two types of curves which have played a significant role in other work on 3D shape representation; namely, *ridges* and *generalized cylinder axes*. A ridge is a reconstruction of an A_3 rib curve and a generalized axis is the A_1^3 axial curve (see Section 3 for notations)³.

2. Related Works

Existing works in 3D \mathcal{MA} computation and regularization have been diverse and differ in the philosophy as well as in the representation. Major \mathcal{MA} computation methods can be broadly organized into five categories based on: (i) *thinning* by peering layers of elements until the \mathcal{MA} is reached, (ii) *ridge following* on the distance map, (iii) solving the PDE of Blum’s *grass-fire* like front propagation [4], (iv) *refining* the Voronoi diagram (\mathcal{VD}) of sample points toward to the \mathcal{MA} , and (v) computing the interior \mathcal{MA} from a solid polyhedral shape, *etc.* For a detailed survey, see [20]. After computing the 3D \mathcal{MA} , the following step of *regularization* is usually necessary (due to the ubiquitous instability) and is in some cases embedded with the prior. Several \mathcal{MA} regularization techniques have been independently developed but rather share common ideas and mainly vary in how the \mathcal{MA} elements are selected to prune. We briefly review them and for in-depth surveys, see [2, 32, 10, 26, 34].

Voronoi refinement: A major approach to 3D \mathcal{MA} extraction is to compute a Voronoi diagram of the sample points and then refine it toward the \mathcal{MA} . In the early works of [24, 3], the interior Delaunay tetrahedra are deleted in layers while maintaining topological consistency. The main problem is that there exist Voronoi vertices near the object surface (centers of flat Delaunay tetrahedra called ‘*slivers*’) that restrain the regularization process. Such ‘*slivers*’ can be filtered out using the *poles* of Amenta *et al.* [1]. In the *PowerCrust* approach [1], the \mathcal{MA} is computed using the *power diagram* (a weighted \mathcal{VD}) of the inner poles. The \mathcal{MA} is then simplified by removing poles of small surface features or with contact balls overlapping significantly. Although this method has some theoretical support, it requires two passes of Voronoi computations and the resulting \mathcal{MA} is not on the Voronoi complex of the input (and needs additional heuristics to clean up). A following work in [10] extracts the \mathcal{MA} by directly filtering the \mathcal{VD} . A Voronoi face is removed if (i) the *angle* between the estimated surface normal and its dual Delaunay edge is not small, (ii) the

³ We point out two related works. (i) The λ -MA [7] gives theoretical support on regularizing the rib curves by pruning all elements closer to the shape boundary (within the threshold λ). (ii) Notably, the approach in [37, 17] extracts smooth rib curves by representing the (Voronoi-refined) \mathcal{MA} as a *two-sided* non-intersecting surface and apply standard Laplacian smoothing (and loop subdivision) on the \mathcal{MA} manifold. However, all interior structural connectivity (junctions) of the \mathcal{MA} is lost.

ratio of the object feature size to its radius is small, which is related to the \mathcal{MA} *significance*⁴. The resulting \mathcal{MA} may contain unwanted internal ‘holes’ and the association between input samples and the \mathcal{MA} can be lost.

Recent development in \mathcal{MA} simplification has focused on retaining the **homotopy** in the process. Informally, the purpose is to ensure that the topology between the shape and the \mathcal{MA} are the same⁵. The work in [33] uses the *PowerCrust* to compute the \mathcal{MA} and ‘*peels*’ off medial sheets according to (i) the sheet size (number of triangles) and (ii) its corresponding shape volume (estimated using the Delaunay tetrahedra), while maintaining the topological consistency. In [31], the *m-Rep* is extracted from an inner Voronoi skeleton via the pruning and merging of medial sheets while minimizing the change of the underlying shape according to two criteria similar to [33]: (i) the sheet area (number of vertices) and (ii) its corresponding shape volume. The recent *flow complex* (\mathcal{FC}) [15] extracts the \mathcal{MA} from the \mathcal{VD} based a *Morse* analysis on the radius flow and is closely related to our approach. A possibly extendable *core* (set of ‘unstable manifolds’) is computed from the flow analysis on the \mathcal{VD} to approximate the \mathcal{MA} . We point out that our approach is based on a *singularity* analysis [12, 20] and produces a richer set than the \mathcal{FC} ⁶.

Refining the \mathcal{MA} of a polyhedron: Another branch of research extracts 3D \mathcal{MA} from a (solid) polyhedral mesh and regularizes it. In [32], \mathcal{MA} branches are pruned using a *separation angle* related to the \mathcal{MA} significance. The homotopy of the \mathcal{MA} is preserved by only removing non-interior medial sheets. In [30, 28] the \mathcal{MA} is computed and pruned simultaneously by measuring the average *flux* of the gradient of the distance field in each voxel.

Curve skeleton: A different branch of research extracts a 1D *curve skeleton* (\mathcal{CS}) out of a 3D shape (instead of the 2D sheets). The \mathcal{CS} is more simplified than the \mathcal{MA} , but its mathematical formulation is still an open question [8]⁷.

3. Representing the \mathcal{MS} Hypergraph Topology

We briefly describe the general topology of the \mathcal{MS} hypergraph, *i.e.*, the *incidence* relationship between the medial sheets (A_1^2), curves (A_3, A_1^3), and nodes ($A_1 A_3, A_1^4$) [12] and propose an efficient data structure to represent it. This representation shall approximate the *true* \mathcal{MA} , captur-

⁴ The *significance* of a \mathcal{MA} branch is related to the corresponding object angle θ and the \mathcal{MA} formation speed v . At any \mathcal{MA} point p , v changes with an angle ϕ between the \mathcal{MA} formation direction and the vector from p to the corresponding object point: $v = -1/\cos \phi$ [11].

⁵ Lieutier shows that any shape are *homotopy equivalent* to its \mathcal{MA} [Lieutier CAD’04].

⁶ All the ‘relay’ type of *shock* vertices (A_1^2 -3, A_1^3 -3) [20] are *not* identified as *critical* points in the Morse analysis in the flow complex framework. Omitting these singular points makes their subsequent analysis and algorithms different from ours.

⁷ The \mathcal{CS} has been related to the \mathcal{MA} via a *geodesic* function on the medial sheets by Dey and Sun [SGP’06].

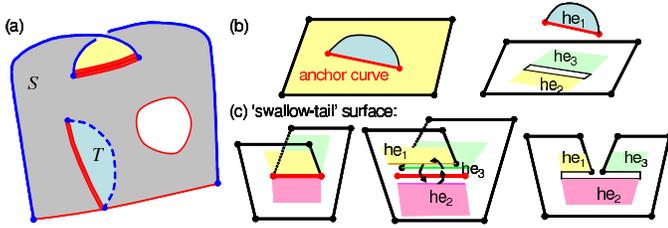


Figure 5. The general topology of an A_1^2 shock sheet \mathcal{S} in the \mathcal{MS} hypergraph and its representation using the *extended half-edge* (\mathcal{EHE}) data structure. (a) Three possible types of sheet-curve incidence of \mathcal{S} : the boundary curve, anchor curve (with a tab T), and swallow-tailed self-intersection. The topology of the 2-incident anchor curve (double red curve) is represented by two half-edges in (b). The 3-incident swallow-tailed self-intersection (triple red curve) is represented by 3 half-edges both in a *loop* (at the junction) and in a *chain* in (c).

ing both global structures and local details, and enable the implementation of hypergraph edits in the transforms. Our solution is a *dual-scale* representation (Figure 4): The *fine-scale* hypergraph is a polygonal mesh consisting of mesh face/edge/vertex elements, typically *non-manifold* at the intersection of medial sheets, which captures detailed information such as the geometry, radius, associated boundary points (*generators*), *dynamics* (speed v , etc.) of the \mathcal{MS} . The *coarse-scale* is a similar structure, but representing groupings of sheets, curves, and nodes towards a *topological* hypergraph also consisting of sheet/curve/node components as hyperlinks/links/vertices, respectively. The coarse-scale representation captures the shape after applying the simplifying transforms.

The grouping of individual sheet elements into a larger scale sheet implies that the coarse-scale medial sheet \mathcal{S} in general may contain the following three kinds of boundaries (Figure 5): (i) a 1-incident *boundary curve* including the A_3 ribs and A_1^3 axials bordering other sheets⁸; (ii) a 2-incident *anchor curve* internal to \mathcal{S} where another medial sheet branch (*tab*) intersects the sheet; (iii) a 3-incident *swallow-tailed* self-intersection of \mathcal{S} close to the A_5 transition (detailed in [13]), which can be viewed as that the A_1^3 curve is triply incident to \mathcal{S} in a loop, and the 3 incidences is also ordered as a single boundary chain (bordering the boundary of \mathcal{S}), Figure 5c.

We adopt the popular *half-edge* (\mathcal{HE}), an edge-centered data structure capable of maintaining incidence information of faces/edges/vertices, to efficiently describe the \mathcal{MS} hypergraph topology⁹. The use of half-edges to describe the incidence relationship between elements of a hypergraph

⁸ The boundary curve of any internal ‘void’ of the sheet \mathcal{S} (which can be either a true hole or part of other sheets) also belongs to this category.

⁹ The original \mathcal{HE} is designed for (and only applies to) a 2-manifold mesh with limitations. It can not represent a mesh with a *degenerate* vertex (which if removed, the mesh is disconnected) or a *degenerate* edge (which is with no incident face). For recent development of 3D *non-manifold* object modeling, see e.g., [Floriani, Hui, SGP’03].

and those of an ordinary mesh is similar. The major difference is that non-planar variations of hypergraph sheets may cause additional topology described above. We extend the original \mathcal{HE} with new capabilities to handle (i) *non-manifold* junctions by extending the half-edge’s *pair* pointer into a circular loop at a junction (Figure 5c), and (ii) *degenerate* mesh configuration by storing the *vertex-link* incidence at each vertex (with a dynamic array). The *extended half-edge* (\mathcal{EHE}) data structure enables all elements in the hypergraph to efficiently traverse all incident objects while keeps only the necessary information. We also adopt a reduced version of the \mathcal{EHE} for the fine-scale \mathcal{MS} mesh, where the topology is more simplified (with planar convex polygons). The coarse-scale \mathcal{MS} components has pointers to access the fine-scale elements as shown in Figure 4. Some details out of scope is omitted due to the space limit.

4. Medial Scaffold (\mathcal{MS}) Transforms

The use of transforms in \mathcal{MA} regularization requires three items: (i) *saliency measure*: estimating the *cost* of corresponding shape deformation, (ii) *detection*: identifying whether any part of a \mathcal{MS} hypergraph is close to any transition and rank order candidate transforms by their saliency measures, and (iii) *operations* of the transform, which includes two parts: the \mathcal{MS} hypergraph edit and the update of the corresponding shape.

The seven transitions of the \mathcal{MA} in Figure 2 cannot all be approached from both sides of the transition. As the red arrows of Figure 2b show, this gives 11 transforms and can be organized into three main categories: *splice*, *contract*, and *merge*. These can be further sub-categorized depending on whether the transform operates on a medial sheet, curve, or node which gives 6 sub-categories: sheet-splice, curve-contract, sheet-contract, node-node merge, node-curve merge, and curve-curve merge. Figure 6 presents a visual summary of the transforms.

The **splice transform** removes a *non-interior* medial sheet (*tab*) T and merges the remaining sheets at the A_1^3 axial curve \mathcal{C} , as illustrated in Figure 6a,b. There are two types of sheet-splice transforms: the A_1A_3 -I prunes and merges sheets as described, and the $A_1^2A_3$ -II contains additional merging of the remaining medial curves \mathcal{C}_1 and \mathcal{C}_2 . The cost of a splice transform is estimated from the tab T ’s saliency using two criteria: (i) the number of associated generators (for small tabs) and (ii) the corresponding change of shape volume divided by the average radius ($\Delta V/r$, for large tabs)¹⁰. With the removal of T , its associated generators are associated with \mathcal{C} for consistency. The change of object shape can be ignored for a small tab; and

¹⁰ Comparing to [33, 31], our use of generators is a better measure for small tabs (assuming rough uniform sampling). To accurately estimate the saliency for larger tabs, incorporating the *local feature size* (w.r.t the radius) is better in keeping small but salient tabs. See [34] for a 2D analysis.

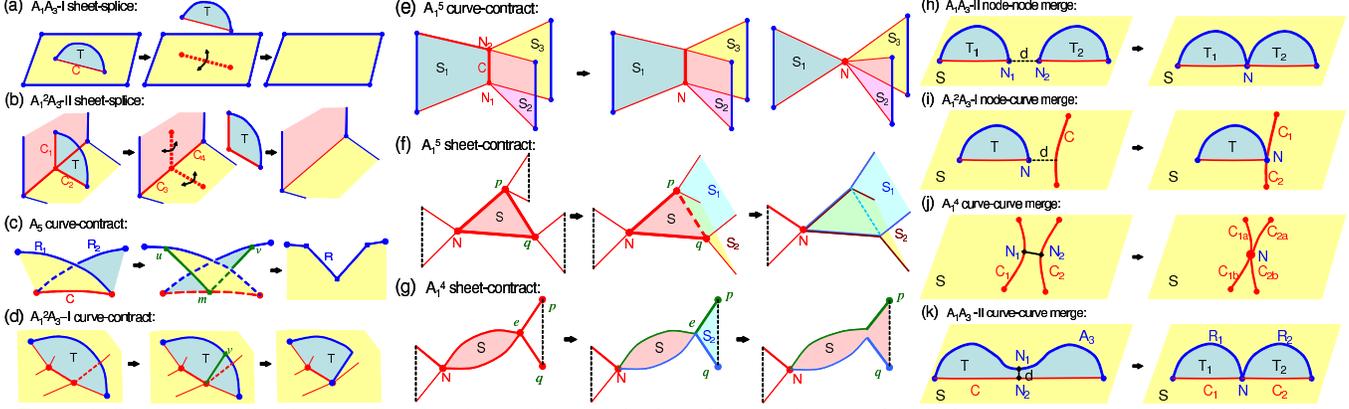


Figure 6. Operations of the eleven types of \mathcal{MS} transforms. (a-b) the *splice* transforms and (c-d) the *contract* transforms on the *boundary* medial sheets. (e-g) the *contract* transforms and (h-k) the *merge* transforms on the *interior* medial sheets of the \mathcal{MS} hypergraph.

for a large tab, a ‘cut-off’ patch of the corresponding protrusion is computed similar to [21] in a zigzag fashion. We observe in practice that the splice transforms simplify the \mathcal{MS} quite significantly, Figures 8 and 9.

The **contract transforms** contract a medial curve \mathcal{C} or a medial sheet \mathcal{S} . Depending on whether the change is near the *boundary* (A_3 rib) or at the *interior* (A_1^3 axial) of the \mathcal{MS} , they are organized into two classes: (i) the A_5 and $A_1^2A_3$ -I curve-contract transforms regularize the boundary, and (ii) the A_1^5 curve-contract and A_1^5/A_1^4 sheet-contract transforms regularize the interior of the \mathcal{MS} . In the A_5 curve-contract transform, the ‘swallow-tail’ of the sheet \mathcal{S} can be removed by identifying a mid-point m on the A_1^3 curve \mathcal{C} and ‘trim’ out the two intersecting wings (by applying splice transforms element-by-element) such that the self-intersection is removed, Figure 6c. The two trimming points u/v on the two rib curves $\mathcal{R}_1/\mathcal{R}_2$ (respectively) are estimated with a proper scale, and the two *trimming paths* um/vm is computed as a *geodesic* shortest path on \mathcal{S} . We approximate the shortest path by applying *Dijkstra’s* algorithm (treating the fine-scale mesh edges/vertices as a graph) so that no cutting through sheet elements is involved. Additional constraints are applied in the shortest path searching to ensure that the swallow-tail topology is completely removed. Similar approach is used in the $A_1^2A_3$ -I curve-contract transform, where only one such trimming is sufficient at the end of the tab \mathcal{T} , Figure 6d. The corresponding shape changes of these two transforms are typically small and can be safely ignored. The costs are approximated using the length of the A_1^3 curve \mathcal{C} under contraction. We observe in practice that the two transforms simplify the topology near the \mathcal{MS} boundary significantly, Figure 9f,g.

The **four** transforms introduced so far simplify the *boundary* of the \mathcal{MS} by removing medial sheets (by elements or by component) and can be performed explicitly on both scales of the \mathcal{MS} . In contrast, the remaining **seven** transforms edit the *interior* of the \mathcal{MS} and require a

‘simulation’ of edits on the coarse-scale hypergraph, while keeping the fine-scale elements intact. Their costs can be roughly estimated using the length of the curve (or area of the sheet) under contraction or merging. The modification of corresponding shape changes require detailed investigation and is omitted here.

In the A_1^5 curve-contract transform, an A_1^3 axial \mathcal{C} is contracted into an A_1^5 node \mathcal{N} by keeping one of \mathcal{C} ’s end point as \mathcal{N} and repeatedly merging \mathcal{C} with other A_1^3 curves on the three incident sheets ($\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3$), Figure 6e. The A_1^5 sheet-contract transform is performed similarly that the sheet \mathcal{S} is separately merged with two sheets \mathcal{S}_1 and \mathcal{S}_2 so that \mathcal{N} is the final A_1^5 node, Figure 6f. The curve pq previously bordering $\mathcal{S}, \mathcal{S}_1$, and \mathcal{S}_2 is virtually removed (in the coarse-scale)¹¹. The above ‘merging’ operations involves *cloning* of fine-scale \mathcal{MS} elements, which can be done by simply adding a pointer to the coarse-scale object. The two scales of representations may thus contain different (but each self-consistent) topology. The A_1^4 sheet-contract transform is performed similarly by merging \mathcal{S} and \mathcal{S}_2 , Figure 6g.

The **merge transforms** merge medial curves or nodes and are also performed in a simulated fashion. In addition, a *geodesic* distance transform is required to detect candidate positions in the \mathcal{MS} hypergraph for these transforms. Both the Dijkstra’s algorithm or the popular *fast marching method* (FMM) can be applied to approximate such a distance field. Figure 7 show two examples in taking all vertex elements on the medial curves as *sources* and computing the distance field on all medial sheets using FMM. In the A_1A_3 -II node-node merge transform, two tabs \mathcal{T}_1 and \mathcal{T}_2 are merged at a node \mathcal{N} , Figure 6h. In the $A_1^2A_3$ -I node-curve merge transform, a node \mathcal{N} is merged with a curve \mathcal{C} and split it into \mathcal{C}_1 and \mathcal{C}_2 , Figure 6i. The A_1^4 and A_1A_3 -II curve-curve merge transforms are similar that two nodes \mathcal{N}_1 and \mathcal{N}_2 are merged into a node \mathcal{N} , Figure 6j,k. We omit implementation details due to the space limit.

¹¹ The A_1^5 node should have 6 incident curves and 9 sheets, the A_1^4 node should have 4 curves and 6 sheets, which are all correctly simulated.

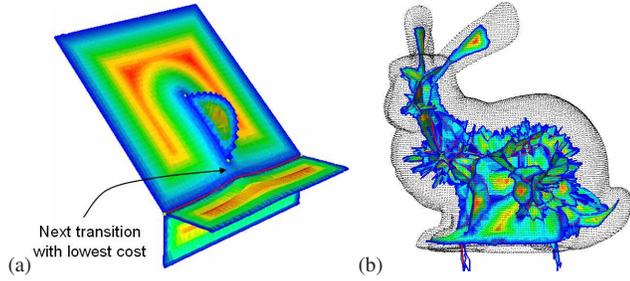


Figure 7. The use of geodesic distance transform on the medial sheets to detect the next possible transition with the lowest cost: (a) The distance map of the \mathcal{MS} near an $A_1^2 A_3$ -I transition taking all vertex elements on the sheet boundary as *sources* (blue). (b) Result of the distance map on the \mathcal{MS} of the Stanford bunny.

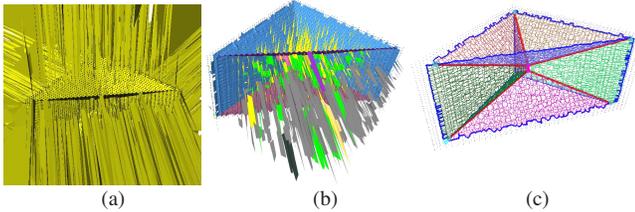


Figure 8. Regularizing the \mathcal{MS} of a prism shape. (a) The immediate \mathcal{MS} after the *shock segregation* (meshing) process [6] contains numerous noisy “spiky” medial tabs. (b) Grouping the medial sheets inside a bounding box into components (yellow: a single sheet element, green: component with 2 elements, gray: sheets connecting to the exterior \mathcal{MS} , other sheets are in random colors). (c) The splice transforms remove the spiky tabs; the remaining \mathcal{MS} hypergraph is shown in wire frame to depict its structure.

The above 11 transforms completely cover all 7 transitions in [13]. Our framework is augmented with two additional transforms, namely, the **gap** and **loop** transforms. The *gap* transform has been used in [6] to segregate the full \mathcal{MS} to reconstruct the surface and yield an initial \mathcal{MS} . The *loop* transform can be used to remove noisy sample points (see [19] for its use in 2D).

5. Computational Scheme for \mathcal{MS} Transforms

The proposed transforms could be applied to regularize the \mathcal{MS} in any arbitrary order, and finding the *optimal* order is computationally expensive and not practical. We consider all transforms in a *greedy* fashion based on their costs. Ideally, we can start with any initial \mathcal{MS} and apply all transforms (including the gaps and loops) until finish. But this is also not practical, since initially the set of medial elements are large and only a few of them shall remain. Our strategy is to (i) group together transforms of the same type with similar costs into iterations, and (ii) perform the transforms that are local and which simplify the \mathcal{MS} effectively, then (iii) construct the coarse-scale \mathcal{MS} and continue to apply all transforms. The computational pipeline is as follows.

Our approach takes unorganized sample points (or polygon clouds with a sampling) as input and compute the *full*

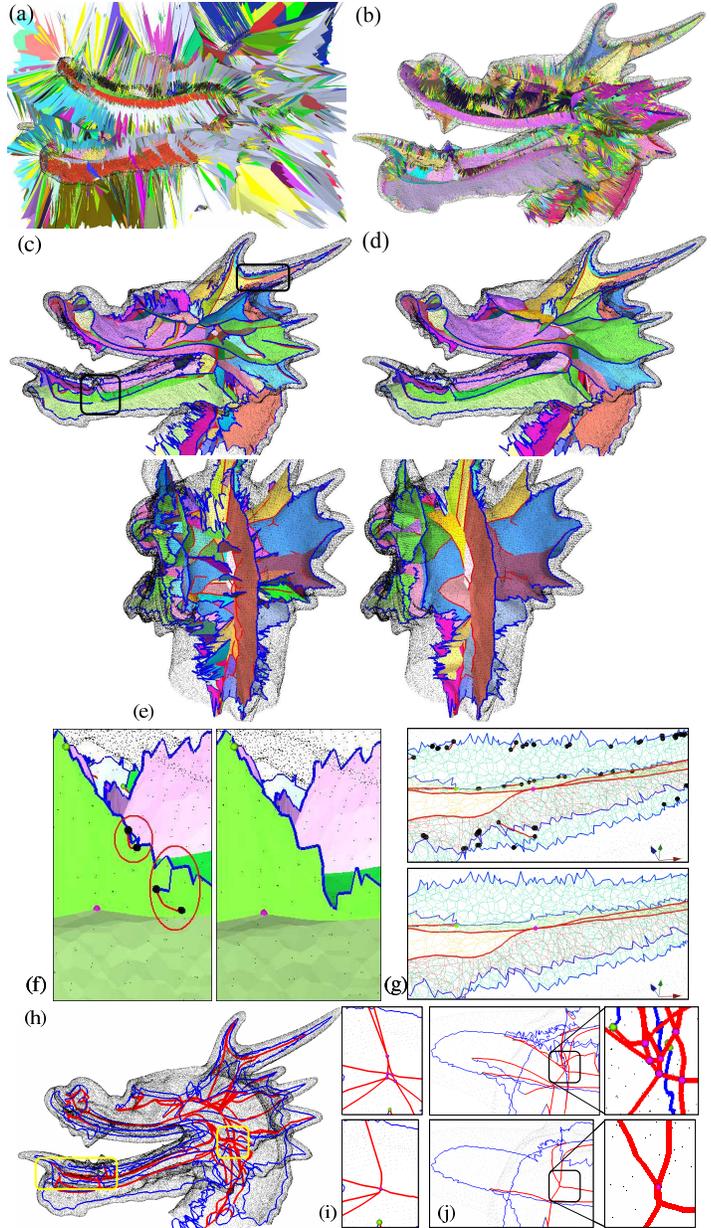


Figure 9. Regularizing the \mathcal{MS} of a complex shape—the Stanford dragon head (127k points). (a) Both the interior and exterior \mathcal{MS} (74k sheets) after *shock segregation* are extremely noisy. The interior \mathcal{MS} (35k sheets) in (b) is simplified in a first stage, *splice regularization*, into 285 sheets, and the coarse-scale hypergraph (1,695 curves, 1,446 nodes) is built as in (c). The \mathcal{MS} hypergraph is then further regularized by a second stage involving all transforms, ending up with only 76 sheets, 262 curves, 219 nodes in (d). (e) shows the effect of *splice* transforms to remove and merge several sheets. (f-g) shows numerous A_5 swallow-tails removed by the *contract* transforms: near the tooth (f) and the horn (g). (h) shows the \mathcal{MS} graph of (d) with sheets implicit to better visualize its interior structure. (i-j) depicts the effect of *contract* and *merge* transforms in simplifying the interior hypergraph topology: near the neck (i) and tongue (j).

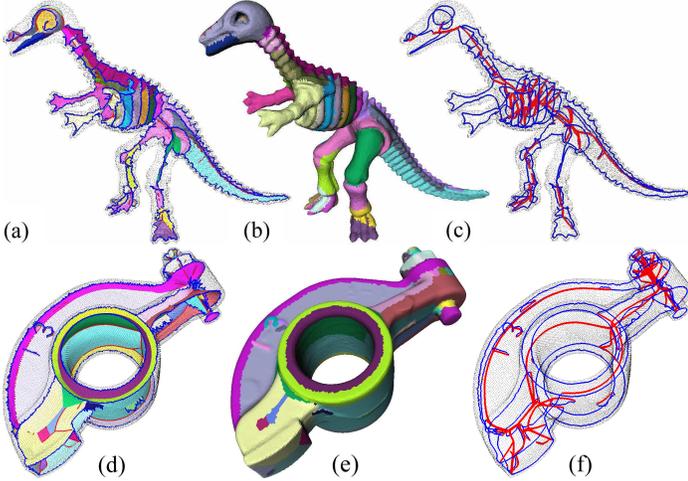


Figure 10. The regularized \mathcal{MS} of (a-c) a toy dinosaur model (14, 050 points) and (d-f) a rocker arm (40, 177 points), data from Cyberware. Both \mathcal{MS} after regularization appear to be neat in structure and are close to the object surface to capture fine details. The tightly-coupled surface regions of the \mathcal{MS} in (b,e) are suitable for further shape modelling and segmentation use.

\mathcal{MS} [20]. The *shock segregation* [6] is then performed, where the gap transforms reconstruct the object surface and produce an initial \mathcal{MS} , which typically contains very noisy ‘spiky’ medial tabs, Figures 8a,b and 9a. We perform the *splice regularization*—a greedy iteration of only splice transforms to remove such small tabs (ordered by their associated number of generators) which greatly simplifies the \mathcal{MS} , Figures 8b-c and 9b-c. The coarse-scale \mathcal{MS} is then built by a connected component analysis. Finally, we apply the greedy iteration considering all transforms on the \mathcal{MS} (the *all-transform regularization*), where all unstable transitions with costs less than a threshold are removed.

Two optional steps can be applied in the above process to refine the \mathcal{MS} : (i) Select the interior/exterior \mathcal{MS} component(s) by *e.g.*, filtering the \mathcal{MS} with a bounding box (Figure 8a-b) and picking the largest \mathcal{MS} component. (ii) Regularize the rib curves in two ways: (a) Apply splice transforms to remove \mathcal{MS} elements with radius smaller than a threshold (an idea related to the λ -MA [7]). (b) Smooth the rib curves using *discrete curve shortening*: move each rib vertex of high curvature toward the bisector of the two neighboring vertices.

6. Experimental Results and Conclusion

We have extensively tested the proposed framework on a large variety of 3D dataset, including artificial shapes to simulate all transitions (Figures 2 and 8), general shapes with salient structure and local details (bunny (Figure 7), dragon (Figure 9), dinosaur (Figure 10a-c), pot fragments (Figure 11)), shapes for industrial applications (fan disk (Figure 4) and rocker arm (Figure 10e-f) from Cyberware), and medical applications (hand (Figure 1), carpal bones

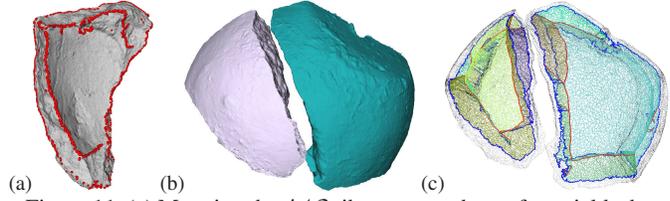


Figure 11. (a) Mapping the \mathcal{MS} rib curves to the surface yields the *ridge* points (in red, see also [17]). (b-c) The structural information of the \mathcal{MS} graph/hypergraph is useful in matching the ‘break curves’ (ridges) in assembling archaeological pot sherds [36].

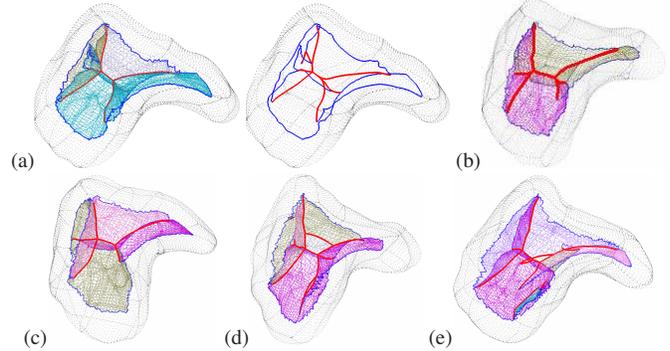


Figure 12. The regularized \mathcal{MS} of the carpal bones (hamates) from several patients are similar in the structure, suggesting the application in shape matching and shape-based diagnosis (data is courtesy of Dr. Crisco, RI Hospital [23]).

(Figure 12)).

Our system handles up to 300k points of input (the bottleneck is the large initial *full* \mathcal{MS} limited by 2GB of computer memory). The regularized \mathcal{MS} is computed from a few seconds to a few minutes. In comparing to existing Voronoi refinement results, ours exhibits three advantages: (i) The \mathcal{MA} is better regularized on both the boundary and the interior topology, (ii) The \mathcal{MS} (rib curves) are closer to the object surface, indicating that our regularization is better in capturing finer details, even in the case of sparsely-sampled or ill-sampled inputs (for example see the dinosaur’s hands in Figure 10a). (iii) All our results are obtained from the sole assumption of unorganized points with reasonable sampling (and nothing further). We can handle shapes with boundary where there is no distinguishing between the interior/exterior \mathcal{MA} .

Applications: The regularized \mathcal{MS} is promising in several shape modeling applications, due to its stability: (i) Figure 11 shows an application to detect *ridges* as well as match the rib curves and the interior structures of the \mathcal{MS} to match the shape fragments. (ii) The corresponding surface regions of the medial sheets are made explicit (Figure 10b,e) to capture the (volumetric) solid of the shape, which also give an initial segmentation of it (which can be further refined). (iii) The regularized \mathcal{MS} is particularly useful in skeleton-based shape matching [5]. Figure 12 shows a set of carpal bones with variations while their \mathcal{MS} are still highly similar in the structure, demonstrating the potential to ro-

bustly match/register 3D shapes. (iv) The regularized \mathcal{MS} also fits well in skeleton-based animation [37] and shape morphing.

Conclusion: We have presented a framework of \mathcal{MS} transforms to stably regularize the \mathcal{MA} across transitions and propose a dual-scale representation to realize it in practice. **Future works** include deriving a better cost estimation and a consistent way to update shape changes for the *interior* contract and merge transforms, so that the \mathcal{MS} transforms can fully model generic shape deformations.

Acknowledgments: This material is based upon work partly supported by the National Science Foundation under Grants NSF/IIS-0413215 and NSF/ITR-0205477.

References

- [1] N. Amenta, S. Choi, and R. K. Kolluri. The power crust. In *Solid Modeling and App.*, pages 249–260, 2001. 3
- [2] D. Attali, J.-D. Boissonat, and H. Edelsbrunner. Stability and computation of the medial axis, in mathematical foundations of scientific visualization, computer graphics, and massive data exploration, 2004. 3
- [3] D. Attali and A. Montanvert. Computing and simplifying 2D and 3D continuous skeletons. *CVIU*, 67(3):261–273, 1997. 3
- [4] H. Blum. Biological shape and visual science. *Journal of Theoretical Biology*, 38:205–287, 1973. 3
- [5] M.-C. Chang, F. Leymarie, and B. Kimia. 3D shape registration using regularized medial scaffolds. In *Proc. 3DPVT*, pages 987–994. IEEE Computer Society, 2004. 1, 7
- [6] M.-C. Chang, F. Leymarie, and B. Kimia. Surface reconstruction from point clouds by transforming the medial scaffold. In *Proc. 3DIM*, pages 13–20, 2007. 1, 2, 6, 7
- [7] F. Chazal and A. Lieutier. The “ λ -medial axis”. *Graph. Models*, 67(4):304–331, 2005. 3, 7
- [8] N. Cornea and P. Min. Curve-skeleton properties, applications, and algorithms. *TVCG*, 13(3):530–548, 2006. 3
- [9] T. Dey, J. Giesen, and S. Goswami. Shape segmentation and matching from noisy point clouds. In *Proc. Eurographics Sympos. Point-Based Graphics*, pages 193–199, 2004. 1
- [10] T. Dey and W. Zhao. Approximating the medial axis from the Voronoi diagram with a convergence guarantee. *Algorithmica*, 38:179–200, 2003. 3
- [11] P. Giblin and B. Kimia. On the intrinsic reconstruction of shape from its symmetries. *PAMI*, 25(7):895–911, 2003. 1, 3
- [12] P. Giblin and B. Kimia. A formal classification of 3D medial axis points and their local geometry. *PAMI*, 26(2):238–251, 2004. 1, 3
- [13] P. Giblin and B. Kimia. Transitions of the 3D medial axis under a one-parameter family of deformations. *PAMI*, accepted, 2007. 1, 2, 4, 6
- [14] P. J. Giblin and B. B. Kimia. Transitions of the 3D medial axis. In *ECCV*, pages 718–724, 2002. 1
- [15] J. Giesen, E. A. Ramos, and B. Sadri. Medial axis approximation and unstable flow complex. In *SCG*, pages 327–336, 2006. 3
- [16] S. Goswami, T. Dey, and C. Bajaj. Identifying flat and tubular regions of a shape by unstable manifolds. In *SPM*, pages 27–37, 2006. 1
- [17] M. Hisada, A. Belyaev, and T. Kunii. A skeleton-based approach for detection of perceptually salient features on polygonal surfaces. *CGF*, 21:689–700, 2002. 1, 3, 7
- [18] N. Iyer, S. Jayanti, K. Lou, Y. Kalyanaraman, and K. Ramani. A multi-scale hierarchical 3D shape representation for similar shape retrieval. In *TMCE*, pages 1117–1118, 2004. 1
- [19] M. Johannes, T. Sebastian, H. Tek, and B. Kimia. Perceptual organization as object recognition divided by two. In *POCV*, pages 41–46, 2001. 6
- [20] F. Leymarie and B. Kimia. The medial scaffold of 3D unorganized point clouds. *PAMI*, 29(2):313–330, 2007. 1, 3, 7
- [21] F. Leymarie, B. Kimia, and P. Giblin. Towards surface regularization via medial axis transitions. In *ICPR*, pages 123–126, 2004. 2, 5
- [22] M. Leyton. *A Generative Theory of Shape*. LNCS 2145. Springer-Verlag, 2001. 1
- [23] D. C. Moore, J. J. Crisco, T. G. Trafton, and E. L. Leventhal. A digital database of wrist bone anatomy and carpal kinematics. *J. Biomech.*, 40(11):2537–42, 2007. 7
- [24] M. Naf, G. Szekely, and *et al.* 3D Voronoi skeletons and their usage for the characterization and recognition of 3D organ shape. *CVIU*, 66(2):147–161, 1997. 3
- [25] S. Pizer and *et al.* Deformable m-Reps for 3D medical image segmentation. *IJCV*, 55(2-3):85–106, 2003. 1
- [26] S. Pizer, K. Siddiqi, G. Székely, J. Damon, and S. Zucker. Multiscale medial loci and their properties. *IJCV*, 55(2-3):155–179, 2003. 3
- [27] T. Sebastian, P. Klein, and B. Kimia. Recognition of shapes by editing their shock graphs. *PAMI*, 26:551–571, 2004. 1
- [28] K. Siddiqi, S. Bouix, A. Tannenbaum, and S. W. Zucker. Hamilton-jacobi skeletons. *IJCV*, 48(3):215–231, 2002. 3
- [29] K. Siddiqi, J. Zhang, D. Macrini, A. Shokoufandeh, S. Bouix, and S. Dickinson. Retrieving articulated 3-D models using medial surfaces. *MVA*, pages 1432–1769, 2007. 1
- [30] S. Stolpner and K. Siddiqi. Revealing significant medial structure in polyhedral meshes. In *3DPVT*, pages 365–372, 2006. 3
- [31] M. Styner, G. Gerig, S. Joshi, and S. Pizer. Automatic and robust computation of 3D medial models incorporating object variability. *IJCV*, 55(2-3):107–122, 2003. 3, 4
- [32] A. Sud, M. Foskey, and D. Manocha. Homotopy-preserving medial axis simplification. In *SPM*, pages 39–50, 2005. 3
- [33] R. Tam and W. Heidrich. Shape simplification based on the medial axis transform. In *VIS*, pages 481–488, 2003. 1, 3, 4
- [34] H. Tek and B. Kimia. Boundary smoothing via symmetry transforms. *JMIV*, 14(3):211–223, May 2001. 1, 3, 4
- [35] N. Trinh and B. Kimia. A symmetry-based generative model for shape. In *Proc. ICCV*, pages 1–8, 2007. 1
- [36] A. Willis, S. Andrews, J. Baker, and *et al.* Assembling virtual pots from 3D measurements of their fragments. In *VAST*, pages 241–254. ACM, 2001. 7
- [37] S. Yoshizawa, A. Belyaev, and H.-P. Seidel. Skeleton-based variational mesh deformations. *CGF*, 26(3):255–264, 2007. 1, 3, 8