

A General Approach to Model Biomedical Data from 3D Unorganised Point Clouds with Medial Scaffolds

Frederic Fol Leymarie^{†1}

Ming-Ching Chang²

Celina Imielinska³

Benjamin B. Kimia⁴

¹Computing Dept., Goldsmiths College, University of London, UK.

²Visualization and Computer Vision Lab., General Electric Global Research Center, Niskayuna, NY, USA.

³Department of Biomedical Informatics, Columbia University, USA.

⁴Division of Engineering, Brown University, USA.

Abstract

We present the latest developments in modeling 3D biomedical data via the Medial Scaffold (MS), a 3D acyclic oriented graph representation of the Medial Axis (MA) [LK07, SP08]. The MS (and associated 3D MA) can be computed as the result of the singularities of a geometric wave propagation simulation. We consider here some of the potential applications of this shape model in the realm of biomedical imaging. We can reconstruct complex object surfaces and make explicit the coarse-scale structures, which are ready-to-use in a number of practical applications, including: morphological measurement for cortex or bone thickness, centerline extraction (curve skeleton) for tracheotomy or colonoscopy, surface partitioning for cortical or anatomical surface classification, as well as registration and matching of shapes of tumors or carpal bones. The MS permits to automatically and efficiently map an unorganised point cloud, i.e., simple 3D coordinates of point samples, to a coherent surface set and associated approximate MA . The derived MS is used to further recover significant medium and large scale features, such as surface ridges and main axial symmetries. The radius field of the MS provides an intuitive definition for morphological measurements, while the graph structure made explicit by the MS is useful for shape registration and matching applications.

Categories and Subject Descriptors (according to ACM CCS): Computer Graphics [I.3.5]: Computational Geometry and Object Modeling—

1. Introduction

Shape understanding and modeling is a central task in biomedical imaging and associated applications. We present the latest development in modeling 3D free-form biomedical data via the *Medial Scaffold* (MS) [LK07, CK08, CLK09], a 3D acyclic oriented graph representation of Harry Blum's Medial Axes (MA), which is closely related to the Voronoi Diagram and its dual the Delaunay Tessellation (in Computational Geometry), Molecular Graphs (in Computational Chemistry) and Critical Nets (in Crystallography) [LK]. The MS (and associated 3D MA) can be computed as the result of the singularities of a geometric wave propagation simulation.

Among the various 3D data input forms and initial shape representations available in biomedical applications, the unorganised point cloud — *i.e.*, an arbitrarily large set of points sampling the surface of objects of interest, with no other information than their associated 3D coordinates — is one of the most common type of input and is the most flexible: objects represented by either volumetric voxels or polygonal meshes can easily be sampled into point clouds. Furthermore, data acquisition such as via scanning, segmentation, or obtained from 3D reconstruction algorithms typically yield geometric information in the form of a set of points.

We present a general approach to model biomedical free-form objects starting from 3D unorganised point clouds by using the MS representation, a hierarchical organisation of the 3D MA into a graph-like structure (Fig. 1). The MS permits to automatically and efficiently map an unorganised

[†] Corresponding Author.




Figure 1: From (a) a point cloud sampling a carpal bone, to (b) its 3D \mathcal{MA} where colors indicate the time flow (distance to point generators), giving a visualisation of local thickness, to (c) the corresponding \mathcal{MS} made of two types of curves, in relation to object surface ridges (in blue) and junctions of \mathcal{MA} sheets (in red).

point cloud to a coherent surface set and associated approximate \mathcal{MA} [LK07, CLK09]. This resulting \mathcal{MA} is used to further recover significant medium and large scale features, such as main surface ridges and axial symmetries. The outcome is ready to use in a range of biomedical applications: (i) morphological measurement (such as bone thickness), (ii) centerline extraction (to model tubular shapes, such as for colonoscopy), (iii) feature detection (of high-curvature regions and ridges), (iv) surface partitioning (for anatomical surface classification), (v) registration for quantification and matching (of molecules, dendrites, bones, tumors).

This paper is organised as follows. We briefly survey related skeleton and symmetry-based 3D shape modeling approaches in §2. We then recall and summarise the definition of the \mathcal{MS} as a generic 3D shape representation in §3. We further discuss various applications of our approach in §4 with an emphasis on surface meshing, shape understanding and shape matching, in the context of cortical thickness measurement, biomedical centerline detection, carpal bone registration. Finally in §5 we outline clinical cases where we propose that the \mathcal{MS} can significantly improve the technical outcome, such as in medical simulation, orthopedic surgery, head trauma modeling, tracheobronchial air-flow modeling, and radiation oncology for tumor shape/volume analysis. We also cover further details on the \mathcal{MS} , especially with regards to its computation, in Appendix A.

2. Medial Representations in Biomedical Imaging

The most common form of medial representation is Harry Blum’s Medial Axis (\mathcal{MA}) first studied in 2D in the 1960’s [Blu73] and in 3D since the 1980’s; while it has been used in many fields, the biomedical field has been a major area of concern from the onset [LK]. For many years, this was restricted to 2D applications, including the shape characterisation of chromosomes, cells, microscopy preparations, slices of tissue from scans [LL92, LL93]. With the advent of increasing processing power, graphics capabilities and sensors modalities, the biomedical field, especially since the 1990’s and with respect to imaging, turned its eye towards the representation of volumetric information, and in particular how to capture the shape of various molecular structures, cells, tissues, bones, objects of interest.

3D medial representations applied to biomedical data can be organised in terms of four main technologies: (i) *Voxel-based*: processing directly a discrete grid of voxels approximating the volumetric shape, (ii) *Geometry-based*: using computational geometry methods to recover intersecting pieces of manifolds (medial sheets bounded by curves ending at nodes), (iii) *Model-based*: predetermined medial primitives are retro-fit to the data, and (iv) *Curve-based*: a simplification of (2D) medial sheets to capture the (1D) main topological features of objects. In what follows we refer to recent works which are specialised for biomedical applications.

Voxel-based. In this category, the input data is usually some pre-segmented volumetric image in the form of voxels representing the object of interest (*aka* the “sugar-cube” model). This volume is then shrunk (thinned) to discard some voxels while selecting only those which approximate a “skeleton” of the interior of the shape. The classical work in this area originates from mathematical morphology and related approaches. The typical algorithm computes a distance map from the boundary and identifies those voxels being at equal distance from that object boundary. Pseudo-Euclidean and Euclidean metrics are commonly used, as well as thinning where the volume is iteratively peeled off preserving only those voxels approximating the \mathcal{MA} [B*, SBS].

Geometry-based. Here the data is usually in the form of a point cloud — a set of point sample coordinates — or pre-segmented in the form of sets of meshed surfaces, such as a triangulation of the boundaries of objects under scrutiny. Computational geometry principles are then applied to compute an exact or approximate \mathcal{MA} . The classical example here takes the form of first computing a Voronoi Diagram (\mathcal{VD}) for the point cloud seen as a set of generators \mathcal{G} of a distance map $L(r, \mathcal{G}) = \min_{G_i \in \mathcal{G}} \|G_i - r\|$. Then, each point generator, G_i , defines a Voronoi region $\mathcal{VR}(G_i)$, the set of points closer to G_i than to any other generators of \mathcal{G} : $\mathcal{VR}(G_i) = \{r \in \mathbb{R}^3, L(r, G_i) \leq L(r, G_j), \forall G_j \in \mathcal{G} - G_i\}$; the Voronoi diagram for m generators is then defined as the collection of all Voronoi regions: $\mathcal{VD} = \{\mathcal{VR}(G_1), \dots, \mathcal{VR}(G_m)\}$. The set of boundaries of these regions is taken as a medial representation of the generators (a graph in 2D and hypergraph in 3D), which is then usually simplified by pruning away branching structures connected to the bounding surface of the sampled object, a result often called *Voronoi skeleton* [Sze, AC].

Model-based. In this category an *a priori* shape model, in the form of primitives like deformable superquadrics or generalised cylinders, is used in a retrofitting process of the input volumetric data, but such that the medial representation of these models is constrained. Pizer *et al.* have used extensively a fixed-topology but deformable (2D) sheet-like medial model for 3D image segmentation [S*03, P*] referred

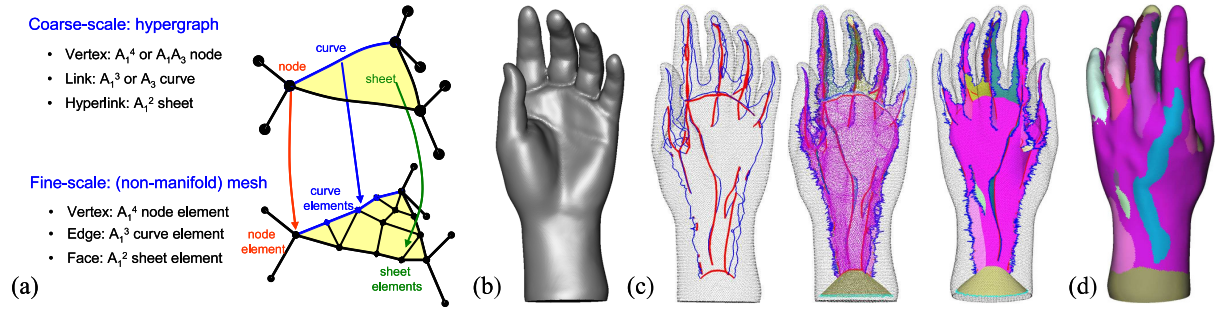


Figure 2: The medial scaffold (MS) as shape representation. (a) The proposed dual-scale MS representation to model the MA of the underlying shape. The coarse-scale MS is a hypergraph capturing the main topology while the fine-scale MS is a mesh capturing the geometric details. (b) Example for a hand point cloud (data from Polhemus Fastscan, 38,219 points) which is automatically meshed into 76,437 triangles using the initial MS for that point set. (c) The regularised MS contains 37 sheets, 131 curves, and 108 vertices (front and back views shown). (d) Some surface regions made explicit from the MS (back hand view). We can directly map medial sheets (loops in the MS graph) to the associated surface regions (generators of the underlying MA symmetries).

to as the *M-rep*. Each such sheet is itself sampled by a few nodes called medial atoms which are connected together in the form of a 2D simple graph structure. Each medial atom is also doubly linked to an approximation of the local bounding surface of the object being represented. An M-rep is thus articulated as each medial atom possesses some degrees of freedom (rotation, elasticity). The initialisation of M-reps can be performed, e.g., by first computing an approximate Voronoi skeleton.

Curve-based. In comparison to the general form of the 3D MA as a set of intersecting (2D) medial sheets, the *curve skeleton* is a connected set of (1D) curvilinear centerlines inside the object, and is used to capture an approximation of that object’s main structures [CM06, T*09, P*09]. Although a curve skeleton provides a much simplified form of medial representation, it cannot capture general shape features, such as surface ridges, and is restricted to objects which resemble connected tubular forms, such as when approximating the human body for animation purpose, or in specialised biomedical applications like for colonoscopy.

A main bottleneck of using the MA to represent and model 3D free-form objects is its instabilities — a slight perturbation of the object’s boundary can add a large but unstable branching structure to the MA. This issue has been recently addressed by representing the MA as the MS and regularising the MS by moving it towards its *transition* points where local topological events occur [GK09].

3. Medial Scaffold (MS) as Shape Representation

The Medial Axis (MA) represents 1st order symmetries between generators of a geometric wave propagation in space — alike light waves which die upon meeting each others: a process also referred to as “grass fire” propagation [Blu73]. In other words, if we consider the classical eikonal (in 3D Euclidean space) for the wavefront $\phi(x,y,z,t)$ [Sta72]:

$\|d\phi\|^2 = 1/c^2$, the MA points are those where wavefronts collapse or first intersect, i.e., where the eikonal first becomes singular — for time collision minima or shortest intersecting geodesic paths. The wavefront propagation itself is simulated using the same point to set distance mapping used to define the \mathcal{VD} , i.e., $L(r, \mathcal{G}) = \min_{G \in \mathcal{G}} \|G - r\|$. Thus, for point generators, like the surface samplings we consider here, the MA symmetries have same trace (geometry) as the boundaries of the \mathcal{VD} . However, the MA adds a time dimension to its trace — i.e., the time of meeting wavefronts or travelled distance to generators — and thus can be thought of as a vector-valued \mathcal{VD} . In 3D the MA takes the form of a complicated set of connected surface patches: the medial symmetry surfaces and their intersection in curves and nodes. There is an equivalence between the time values of meeting wavefronts and the radii of maximal balls touching the surface at sample points, such that the centres of these balls precisely trace MA points. We use the A_k^n notation to refer to such maximal balls with k -th order of contact at n surface (or sample) points [GK04]. A_3 and A_1^3 shock curves delimit A_1^2 MA sheets; for a point cloud, A_1^3 curves are identical to Voronoi edges, their A_1^4 endpoints are Voronoi vertices, and their duals are Delaunay triangles through the $n=3$ sample/contact points. A_1^3 curves are computed from a sequence of critical points of the radius flow (the radius of maximal balls projected on the MA): sources for MA sheets (A_1^2 -2) are paired to find sources for shock curves (A_1^3 -2), which in turn are paired to find A_1^4 endpoints. An in-depth analysis can be found in the paper by Leymarie and Kimia which presents a “geometry-based” set of results and implementation [LK07]. A “voxel-based” implementation has also been studied by Leymarie in his PhD thesis [Ley03, Ch.5].

For any given object input, the MS is a simplified oriented 1D graph structure defined to represent the 3D MA in its most compact form, while preserving its uniqueness.




Figure 3: *MS for modeling bone shapes. (a) The re-meshed surface of a hip bone (132,538 points, 264,927 triangles, point cloud data from Cyberware); the regularised (interior) MS (299 sheets, 475 curves, and 349 vertices), sheets colored by radius (distance field), which intuitively suggests the thickness measure; axial curves in red, ribs of ridges in blue) and corresponding surface regions. (b) Mandible data (35,555 points meshed into 71,105 triangles); MS sheets colored by radius field; MS axial and rib curves, MS sheet components, and the ridge points (red) on the shape surface. (c) Vertebral data (9,589 points meshed into 19,193 triangles); MS sheets colored by radius field and by components; surface regions correspond to MS sheet components; and the detected ridge points. Observe how the salient features of the bone shapes are well captured in the regularised MS, which is useful for further diagnosis applications.*




Figure 4: *MS for modeling the cortex shape (from within). (a) Brain (WM) point cloud (110,797 points) meshed into 220,477 triangles. (b) MS sheets colored by radius (distance) field, where the dark blue regions correspond to high curvature folding. (c) MS sheets colored by individual components, where corresponding surface boundary components are shown in (d).*

The *MS* explicitly represents only the singularities of the time (or radius) function along the *MA* trace together with boundary conditions and junctions taking the form of a set of special nodes connected by special curves (Fig. 2). These singularities — the sources, relays and sinks of the time flow along the *MA* — and their curvilinear connectivity paths — as indicated by the time flow — are taken as the special nodes and curves building the acyclic graph structure we call the *MS* [LK07]. The *MS* preserves the nice properties of the *MA*, and thus provides for a compact representation of

shape itself. Specifically the *MS* graph structure can be effectively regularised by applying a set of transforms (graph edit operations) to simplify it toward *MA* transition points [GK09,CK08]. These are higher order symmetry points simulating a change of topology of the *MA* under a continuous deformation. This provides a formal framework to deal with varying scales of sudden changes in topology, from perturbations due to noise to those due to articulated motions. The *MS* is also a natural structure to conduct shape characterisation and recognition tasks [CK09]. In Appendix sec:app we

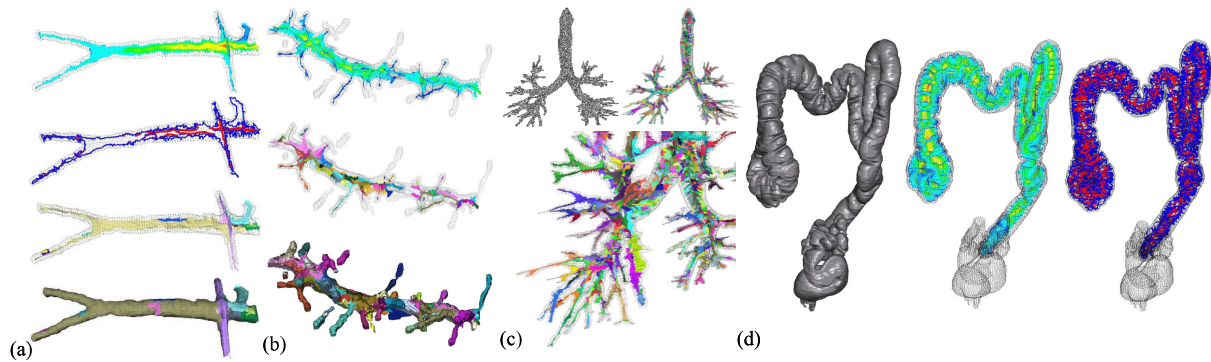


Figure 5: *MS* for modeling tubular and axial biomedical shapes. (a) Aorta (3,843 points) meshed into 7,651 triangles, where the *MS* sheets are colored by radius (distance) field; *MS* axial curves in red and rib curves in blue; and the *MS* sheets and corresponding shape surfaces are colored by components. (b) Dendrite (12,478 points) meshed into 18,025 triangles; the *MS* sheets and corresponding surface components. (c) Bronchial tree (84,615 points and 168,918 triangles) and its *MS*. Observe how close the medial sheets approaching the object boundary capture the details of shape. (d) Colon (104,392 points) meshed into 186,538 triangles and its *MS*. The axial curves after regularisation lead to an automatically computed centerline useful for virtual colonoscopy.




Figure 6: *MS* for detecting salient surface ridges. (a) The *MS* sheet boundary at an A_3 rib curve maps to the ridge on the object surface. (b) Local configuration of mapping the “ridge vector” from an A_3 rib to the ridge (see text). (c) Result of ridge detection on a face (11,748 scan points, data from MPII). Ridge points are in red and the two ridge region curves are in cyan and yellow, respectively. Observe that only coarse-scale “convex” ridges (from within the face mask) such as the nose are detected, while other fine-scale features are suppressed. A similar analysis from outside leads to the detection of coarse scale valleys.

provide more details on the computation and regularisation of the *MS*.

4. *MS* for Biomedical Shape Modeling

We discuss a wide extent of applications by representing biomedical and anatomical shapes using the *MS*, ranging from surface meshing of input point clouds, shape understanding via extracting salient axial symmetries and high curvature surface ridges, and object registration by matching the corresponding compact *MS* graph structures.

4.1. Surface meshing for morphological measurements

Meshing input point cloud into surface meshes. The *MS* is typically first applied to the raw input point cloud to automatically retrieve a useful surface set. The part of the computed *MS* which maps to local symmetries of points near

the object’s surface is used to obtain a Delaunay triangulation of that surface; this method makes minimal assumptions on the data, and can deal with surfaces which may be non-oriented, non-continuous (with sharp surface features), which may have holes and boundaries, or which are smooth and bounding solids, in the same framework [CLK09], making it particularly suited for biomedical imaging where the surfaces encountered can have various topologies and geometries. The remaining parts of the *MS* — not used to reconstruct a surface mesh — then provide for a useful approximate *MA* of the underlying object.

Morphological Measurement. The radius field of the *MS* makes explicit the radial dimensions of solid shapes, thus providing a generic and efficient way to quantify the thickness of objects. Fig. 3 (2nd column) shows examples of




Figure 7: The MS (and associated MA) was computed for cortical thickness measurement. The input MR brain scan, from the OASIS dataset, is first segmented via a method by Chang and Tao [CT10]. (a) The gray matter (GM), of the cerebral cortex, bounded by two surfaces (namely, the “inner surface” (IS) and the “pial surface” [CT10]) are then extracted. (b) The MA of the GM permits to recover well the trace of the ridges of the GM, thus providing a consistent and robust approach to measure cortical thickness. Here the MA is colored by the distance to its two boundary surfaces, which is half of the GM thickness.

bone thickness, while Figures 4 and 7 show cortical thickness measurements.

4.2. Making explicit salient axial structures

Detection of salient axial symmetries of shape. The A_1^3 axial curves of the MS (red curves in Figures 2, 3, 5) correspond to the axial symmetries of the shape, which is the trace where three (or more) medial sheets intersect.

Centerline detection. Further simplification and reduction of the MS , e.g. by shrinking medial sheets following a medial geodesic function [DS06], lead to a simplified 1D form of a curve skeleton (Fig. 5). This is useful in medical application such as human tracheobronchial tree modeling and path planning in virtual colonoscopy.

4.3. Making explicit salient surface structures

Ridge detection. Ridges are important salient features of a free-form object where its surface bends sharply. Mathematically, a ridge curve is made of surface points where the magnitude of the largest principal curvature attains an extremum along its corresponding lines of curvature. The MS makes explicit surface ridges, in that medial sheet outer boundaries (A_3 ribs, blue curves in Figures 2, 3 and 5) directly correspond to such ridges. Specifically, a set of ridge points can be computed by mapping an A_3 rib curve to the corresponding ridge curve on the surface. Such mapping is guided by a “ridge vector” (red arrows in Fig. 6(b)), which can be computed from a bordering medial sheet element (polygonal patch S) toward each corresponding ridge point on the surface: each such patch S can be associated with two boundary sample points (G_a, G_b), which together with S define a fan-like region perpendicular to S . The collection of such fans is then used to trace by interpolation the surface ridge. In Fig. 6(c), many salient convex face features such as the nose, chin, eyebrows, cheeks, and lips are made explicit by the ridges. Fig. 3(b,c) illustrates convex ridge detection on bone

shapes (mandible and vertebral). When applied “outside” the object’s surface, the same ridge analysis maps valley lines of concavities.

Shape Partitioning. The MS sheet components make explicit a partition of the underlying object on the basis of its MA symmetries. Fig. 4 gives an example for sulci and gyrus classification and Fig. 5.(a and b) give partitions of the branching components of an aorta and a dendrite.

4.4. Shape matching and registration

Registration and matching of anatomical shapes. The compact graph structure of the MS provides a simple and useful representation to perform shape matching and registration computations. Fig. 8 illustrates an example of matching carpal bones. This is useful for producing a computational atlas of shapes. In particular, this shape matching capacity is useful in the quantification and registration of tumors, permitting to accurately determine their volumetric behavior, as will be discussed in the next section.

5. Case Discussions: The MS in Medical Applications

We now discuss the potential of the MS in a variety of clinical/medical applications. We outline each application case-by-case and point to the benefits of using the MS .

Medical simulation – mandible. In medical simulation, the key training application is the development of dynamic anatomy visualisation for teaching, computer-assisted surgery (CAS), and treatment planning [A*03, IM05]. Usage of patient-derived anatomy and physiology makes the simulation relevant to patient specific planning and clinical intervention. For example, modeling human mandible from patient data and superimposing it on a human patient simulator (HPS), requires matching two sets of landmarks on a virtual mandible and a real target-specific mandible. We argue that landmarks on mandible and related shape characteristics, such as geodesic distances between landmarks and

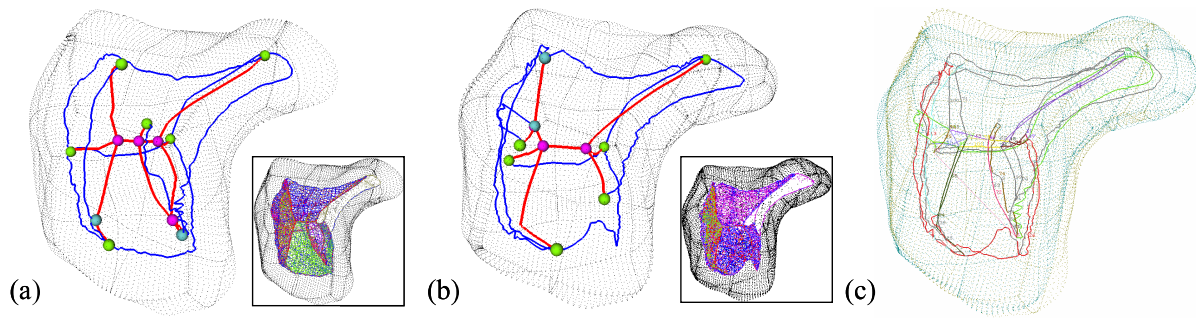


Figure 8: The MS of two carpal bones (Courtesy of Dr. Crisco, RI Hospital) and the matching results in (c) [CK09].

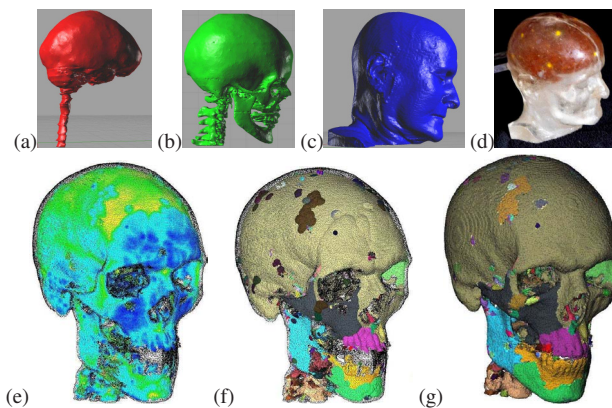


Figure 9: From raw CT data to virtual (a-c) head model and physical (d) head phantom. (e-h) A simulation using the MS on a dataset of 272,171 points meshed into 462,521 triangles. (e) The MS radius field which reflects bone thickness. (f) The MS with 317 sheets and corresponding surface partitioning in (g).

ratios of the measured distances, can easily be expressed using the MS (e.g., Fig. 3(b), and recent results (e.g., [Hea08]) could be improved and refined with this new approach.

Orthopedic surgery. The maturing area of designing custom-made prosthesis for hip joint revision arthroplasty nowadays employs Computer-Aided Design and Manufacturing technologies (CAD-CAM) [D*07]. This requires prosthetic devices that are designed and manufactured with high precision to match anatomical and biomechanical characteristics of a patient. The quality of the custom-made prosthesis depends on the accuracy of the segmented relevant structures: pelvis and femur. An error of segmentation, a crucial component on the imaging, visualisation, and modeling pipeline, may greatly impact the final prosthesis design. Similarly, the quality of modeling of the segmented pelvis and femur that follows for the virtual simulation of the reconstructed joint, and manufacturing process for the realisation of a prosthesis, depends on the ability to derive the correct curvature and thickness of the bones, and to quantify shape characteristics. This requires storing information

about shape and landmarks which can be made explicit by the MS (e.g., Fig. 3(a), thus offering ways to improve existing processes of hip joint custom-made design of prostheses in CAD-CAM systems.

Modeling head trauma under blast. Anatomy and medical image based high-fidelity computational modeling can be used to analyse trauma injury to the brain, lung, spinal cord, gastrointestinal track, due to blast explosion and impact in both civilian and military population [IPT06]. Physical head phantom that include, skin, skull, CSF and brain (with pons and ventricles), built with materials that best mimic biomechanical properties of the actual tissue of these head structures, are used to collect time-varying data for head trauma under simulated conditions to study traumatic brain injuries (TBI). This type of novel approach for the testing of virtual models in simulating head trauma resolved by FEM, which is based on creating a realistic head phantom used to collect physical evidence, provides unique insights into the understanding of the complexity of head injuries. In Fig. 9, we depict the results of the following prototyping steps: (i) 3D segmented and reconstructed, from a CT scan (in this case) models of brain, skull, and facial features on the skin; (ii) models are converted into CAD format files; (iii) models are verified with 3D measurements and fit; (iv) rapid prototype patterns for head structures are created. This process can be greatly improved when using the MS to model the prototyping steps (Fig. 9(e-g)). Hence, we could have better control of capturing landmarks, shapes, required thickness of the modeled structures, all features which are made explicit by the MS .

Modelling air-flow in the tracheobronchial tree. Simulation of the air-flow inside the tracheobronchial tree can provide patient specific modeling of air motion inside the lungs [K*07]. The geometry, shape, length and dimensions of anatomical structures vary among individuals, different age groups, and gender [GW86]. It is crucial to extract a correct computational model of the geometry and topology of the tracheobronchial tree for the biomechanical model to work correctly. Accurately segmented patient-specific models, obtained from 4D CT data, could also be represented by the MS and made ready for modeling mechanical properties

of the pulmonary system. The \mathcal{MS} can provide more flexible, elegant representations that will allow quantification and tracing of features. This simulation framework can be applicable for lung radiation therapy planning, tumor motion prediction, and assessment of malfunction of lung tissue.

Quantification of tumor volume/shape for image-guided radiation therapy. In radiotherapy, segmentation is used to derive tumor volume and to compute radiation dose to eradicate the tumor and protect the normal tissues surrounding the tumor. There is a lack of standardised approach to automate the segmentation of a tumor, evaluate its volume, and quantify the treatment response. There are various available segmentation techniques, such as those provided by the National Library of Medicine Insight Segmentation and Registration Toolkit (ITK) (www.itk.org). In a study, a variety of segmentation methods should be tested, including hybrid [I*04] and skeletal [TK09] ones where each tumor will be delineated using a number of algorithms with a range of parameter settings. The \mathcal{MS} can then be applied to quantify the differences in volume and selected shape features, and the results can further be validated against ground truth, where volume and shape are derived from a surgical pathology specimen.

In the variety of applications, as we described in selected examples, we argue that the \mathcal{MS} is an elegant and generic tool, capable of playing a similar role as once the marching cubes algorithm did [LC87], in modeling the surfaces of complex shapes in biomedicine. In addition, the \mathcal{MS} provides means to quantify a variety of shape features, thus combining modeling and feature quantification in one unified approach.

6. Conclusion

We have reported on our progress in using and refining the classical medial representation of shape — Blum's Medial Axis (\mathcal{MA}) — in the form of a compact 3D graph structure: the Medial Scaffold (\mathcal{MS}), recently introduced in Computing [Ley03, LK07, Cha09, CLK09]. This representation of shape permits to efficiently map large unorganised point clouds to useful surface meshes at a fine scale, and make explicit a number of salient shape structures at larger scales, including surface ridges, surface and volumetric parts, local shape thickness. We have also illustrated biomedical applications of this shape representation framework, including examples of morphological measurements, shape analysis and registration. The next step for this work is to further demonstrate, and improve where needed, the technical outcomes in specific biomedical applications, such as combined accuracy and efficiency control in clinical scenarios.

References

- [A*03] ARGOTTI Y., ET AL.: Dynamic superimposition of synthetic objects on rigid and simple-deformable real objects. *Computers & Graphics* 26 (2003), 919–930.
- [AC] AMENTA N., CHOI S.: *Voronoi Methods for 3D Medial Axis Approximation*. ch. 7, pp. 223–240. in [SP08].
- [B*] BORGEFORS G., ET AL.: *Discrete Skeletons from Distance Transforms in 2D and 3D*. ch. 5. in [SP08].
- [Blu73] BLUM H.: Biological shape and visual science. *Journal of Theoretical Biology* 38 (1973), 205–287.
- [Cha09] CHANG M.: *The Medial Scaffold for 3D Shape Modeling and Recognition*. PhD thesis, Brown, 2009.
- [CK08] CHANG M., KIMIA B.: Regularizing 3D medial axis using medial scaffold transforms. In *IEEE Proc. CVPR* (2008), pp. 1–8.
- [CK09] CHANG M., KIMIA B.: Measuring 3D shape similarity by matching the medial scaffolds. In *IEEE 3DIM* (October 2009).
- [CLK09] CHANG M., LEYMARIE F., KIMIA B.: Surface reconstruction from point clouds by transforming the medial scaffold. *CVIU* 113, 11 (2009), 1130–1146.
- [CM06] CORNEA N. D., MIN P.: Curve-skeleton properties, applications, and algorithms. *IEEE TVCG* 13, 3 (2006), 530–548.
- [CT10] CHANG M., TAO X.: Subvoxel segmentation and representation of brain cortex. In *SPIE Medical Imaging: Image Processing* (Feb. 2010), vol. 7623.
- [D*07] DOMANSKI J., ET AL.: Evaluation of automated segmentation of hip joint in revision arthroplasty. *J. CARS* 2 (2007), 246–272.
- [DS06] DEY T., SUN J.: Defining and computing curve-skeletons with medial geodesic function. In *SGP* (2006), Eurographics, pp. 143–152.
- [GK04] GIBLIN P., KIMIA B.: A formal classification of 3D medial axis points and their local geometry. *IEEE Trans. PAMI* 26, 2 (2004), 238–251.
- [GK09] GIBLIN P., KIMIA B.: Transitions of the 3D medial axis under a one-parameter family of deformations. *IEEE Trans. PAMI* 31 (2009), 900–918.
- [GW86] GRISCOM N., WOHL M.: Dimensions of the growing trachea related to age and gender. *American Journal of Roentgenology* (1986), 233–237.
- [Hea08] HIPPALGAONKARN., ET AL.: Generating classes of 3D virtual mandibles for AR-based medical simulation. *Simulation in Healthcare* 3 (2008), 103–110.
- [I*04] IMIELINSKA C., ET AL.: Hybrid segmentation methods. In *Insight Into Images*. A.K. Peters, 2004, ch. 12, pp. 351–388.
- [IM05] IMIELINSKA C., MOLHOLT P.: Incorporating 3D virtual anatomy into the medical curriculum. *Commun. ACM* 48, 2 (2005), 49–54.
- [IPT06] IMIELINSKA C., PRZEKWA A., TAN X.: Multiscale visual analysis of trauma injury. *J. of Information Visualization* 5 (2006), 279–289.
- [K*07] KAYA I., ET AL.: Modeling air-flow in the tracheo-bronchial tree using computational fluid dynamics. In *Comput. Biomechanics for Medicine Workshop* (2007).
- [LC87] LORENSEN W. E., CLINE H. E.: Marching cubes: A high resolution 3D surface construction algorithm. *Computer Graphics* 21, 4 (1987).
- [Ley03] LEYMARIE F. F.: *3D Shape Representation via Shock Flows*. PhD thesis, Brown, 2003.
- [LK] LEYMARIE F. F., KIMIA B. B.: *From the Infinitely Large to the Infinitely Small*. ch. 11. in [SP08].
- [LK07] LEYMARIE F., KIMIA B.: The medial scaffold of 3D unorganized point clouds. *IEEE Trans. PAMI* 29, 2 (February 2007), 313–330.

- [LL92] LEYMARIE F., LEVINE M.: Simulating the grassfire transform using an active contour model. *IEEE Trans. PAMI* 14, 1 (Jan. 1992), 56–75.
- [LL93] LEYMARIE F., LEVINE M.: Tracking deformable objects in the plane using an active contour model. *IEEE Trans. PAMI* 15, 6 (June 1993), 617–634.
- [P*] PIZER S., ET AL.: *Synthesis, Deformation, and Statistics of 3D Objects via M-Reps*. ch. 8. in [SP08].
- [P*09] PATANÈ G., ET AL.: A minimal contouring approach to the computation of the Reeb graph. *IEEE Trans. VCG* 15, 4 (2009), 583–595.
- [S*03] STYNER M., ET AL.: Automatic and robust computation of 3D medial models incorporating object variability. *IJCV* 55, 2-3 (2003), 107–122.
- [SBS] SIDDIQI K., BOUIX S., SHAH J.: *Skeletons via Shocks of Boundary Evolution*. ch. 4. in [SP08].
- [SP08] SIDDIQI K., PIZER S. M. (Eds.): *Medial Representations: Mathematics, Algorithms & Applications*, vol. 37 of *Comp. Imaging & Vision series*. Springer, 2008.
- [Sta72] STAVROUDIS O. N.: *The Optics of Rays, Wavefronts, and Caustics*, vol. 38 of *Pure and Applied Physics*. Academic Press, 1972.
- [Sze] SZEKELY G.: *Voronoi Skeletons*. ch. 6. in [SP08].
- [T*09] TAGLIASACCHI A., ET AL.: Curve skeleton extraction from incomplete point cloud. In *ACM Transactions on Graphics* (2009), vol. 28.
- [TK09] TRINH N., KIMIA B.: Category-specific object recognition and segmentation using a skeletal shape model. In *BMVC* (2009).

Appendix A: Medial Axis and Scaffold

In this appendix we summarise the important properties of the \mathcal{MS} many of which it inherits from the \mathcal{MA} , and we also summarise how the \mathcal{MS} is computed for the applications presented in this paper.

The \mathcal{MA} is a generic representation of shape which is used to study the highly regular and inorganic objects, for example found in crystals or crafted by humans and machines, as well as the irregular, more free-form and organic objects found in biomedical applications, such as bodies, tissues, leaves, arteries, cells, macro-molecules [LK]. The principal properties of the \mathcal{MA} which makes it a good candidate as a generic shape representation can be listed as follows [LK07]: (a) it is intuitive in representing elongated and branching objects; (b) it encodes the varying width of forms via the radius function associated to each \mathcal{MA} point; (c) useful boundary features are made explicit by the \mathcal{MA} extremities, including curvature extrema, corners, ridges, valleys; (d) other geometric and topological features are made explicit, including holes, necks, branching and looping sub-structures; (e) it provides a hierarchy of scales of sub-parts via the (\mathcal{MS}) branching graph structures; (f) it is *complete* in the sense that, before pruning its graph sub-structures or regularising it, the original object trace is always recoverable by a reverse wave propagation initiated at the \mathcal{MA} points and limited by




Figure 10: From (a) a point cloud sampling a prism, to (b) its 3D \mathcal{MA} where colors indicate the radius flow (distance to point generators), giving a visualisation of local thickness, to (c) the corresponding \mathcal{MS} made of two types of curves, in relation to object ridges (in blue) and junctions of \mathcal{MA} sheets (in red). (Adapted from [Ley03], Fig. 2.1.)

the associated radius values; (g) shape dynamics (perturbations, deformations, growth and kinematics) are made explicit via the graph layout and associated radius flow.

In 3D applications, the \mathcal{MA} suffers from being a complicated set of intersecting surfaces (A_1^2 sheets). The \mathcal{MS} provides a formal and efficient way to represent the \mathcal{MA} as a graph made of 1D structures only built from medial curves of two main types: axial curves (A_1^3 always shown in red in our figures) where A_1^2 sheets intersect and rib curves (A_3 always shown in blue in our figures) which delimit surface ridges. The \mathcal{MS} is built from special points along these two types of curves where the radius function goes through extrema: these are the singularities of the radius flow along the \mathcal{MA} curves. Because of the looping topology of the \mathcal{MA} , where any A_1^2 sheet is necessarily bounded by either (i) sets of A_1^3 curves (connected via A_1^4 endpoints) or (ii) sets of A_1^3 and A_3 curves (connected via A_1^4 and A_1A_3 endpoints), the \mathcal{MA} can be recovered from the \mathcal{MS} , where the missing information about the interior of the A_1^2 sheets can be regenerated from the radius flow along each of the \mathcal{MS} graph loops (identically bounding an original A_1^2 sheet) which provide boundary conditions for each \mathcal{MA} sheet [LK07]. Thus in 3D the \mathcal{MS} offers an ideal way to “compress” the redundant information still present in the \mathcal{MA} , while making available the graph structure essential in applications for the study of shape (Fig. 10).

In terms of computations, as we elaborated in §1 and §3, the most practical way is to start from a point cloud which is assumed to sample a pre-segmented object. Each initial point sample is seen as a generator G_i of a 3D Euclidean (spherical) wave propagation. As wavefronts propagate, the initial \mathcal{MA} contacts are found for closest pairs of generators, the sources of medial sheets, denoted A_1^2-2 . These initial sources for the construction of the \mathcal{MA} are efficiently computed as mid-points of pairs of generators defining a maximal ball (*i.e.*, not including any other generator). Leymarie and Kimia have showed how such computed \mathcal{MA} sheet sources can also be paired to directly compute the initial intersection of the corresponding \mathcal{MA} sheets: the sources of \mathcal{MA} axial curves denoted A_1^3-2 [LK07]. Each such pair of A_1^2-2 's provides a unique triplet of generators (3 initial point samples) which together can be used to compute a candidate A_1^3 source of an axial curve at the barycenter of the




Figure 11: This figure illustrates how a MS curve (at the junction of 3 MA sheets) is associated to a surface triangle in a deformation simulating a reverse sampling process [CLK09]. Each time pairs of point samples (generators of a wave propagation) collide, an MA sheet can be associated to a (Delaunay) edge. Once 3 associated such pairs have been identified, what only remains is the corresponding MA curve and its associated (Delaunay) triangle.

corresponding triangle — taking as triangle vertices the 3 generators. Again, if this candidate point is at the locus of a maximal ball (not including any other generators) it remains as a valid MA (and MS) point (Fig. 11). Different configurations of barycenters (with corresponding triangles being acute or obtuse) then lead to an iterative pairing of A_1^3 -2 sources of axial curves to define unique quadruplets of generators from which candidate axial curve endpoints are obtained at their respective barycenters, each made of 4 generators; the corresponding tetrahedra being of three possible types (as a function of the location of the barycenter in or out the tetrahedron). Each such candidate which is at the center of a maximal ball defines an MA (and corresponding MS) vertex denoted A_1^4 . An iterative process of pairing A_1^3 axial curves then takes place as a function of the different possible configurations of triangles and tetrahedra of generators. The process is guaranteed to terminate with the exact and full set of special MA points needed to construct the MS . The latter is defined as a connected and oriented graph set by the linking curves through all A_1^3 -2 sources and A_1^4 endpoints; the A_1^3 -2 sources of MA sheets, now redundant, can then be discarded.

The above process defines the initial step in computing and using the MS in practical applications. The second step consists of using a subset of all A_1^3 axial curves to construct a polygonal approximation of the surface bounding the pre-segmented object. This can either be done (i) without relying on the initially segmented bounding object surface (in the case where only point samples were reliably obtained) [CLK09] or (ii) with knowledge of the true object surface: in either case, any A_1^3 axial curve which crosses the bounding surface neighborhood can be separated from the rest of the MS . This “segregation” process reliably produces two sub-structures for the MS : (a) one in association to a triangulation of the bounding surface where each individual A_1^3 curve is mapped to a surface triangle, and (b) a graph approximation of the MA of the segmented object (Fig. 12).

Note that the resulting MS approximates the MA inside as well as outside the object; in all our examples in this paper we only illustrate the interior MS (or corresponding interior

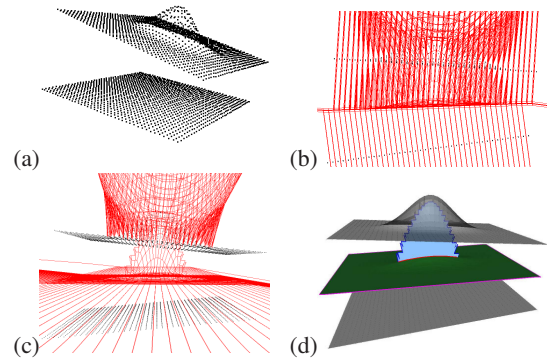


Figure 12: Illustration of the MS segregation process (after [Ley03, Ch.6]). (a) A set of 3,200 points are uniformly distributed on a pair of planes, one of which is deformed by an elongated Gaussian kernel; (b) side-view of the corresponding full MS ; (c) the MS sub-parts not crossing the object's surface neighborhood; (d) the results of this shock segregation are two-fold: (i) the reconstructed surface and (ii) its corresponding MS hypergraph.

MA). Both the interior and exterior MS sub-parts are useful in practice: for example in the cortical study, the interior MS makes explicit convex regions and surface ridges while the exterior MS makes explicit concavities, deep folds and valleys. When holes exist in the segmented object, the “exterior” and “interior” MS will remain connected after segregation, via sub-structures corresponding to each hole; these can in turn be used to repair the object's segmentation for further use and analysis, e.g. to improve registration results.

The third step is to structurally simplify the MS in a step by step fashion by *regularisation* [CK08, Cha09]. At each iteration, a transform which simplifies the MS topology most, and modifies the corresponding shape least, is selected. The iterative regularisation stops when a predefined criteria is reached, such that the application of further transforms becomes prohibitive, for example by requiring to modify the object's shape by too large an amount.

All our examples in this paper start from unorganised point clouds which are processed via these three steps: (i) full MS computation [LK07], (ii) segregation of the parts directly crossing the object surface [CLK09], (iii) simplification by regularisation [CK08]. These computations are either automatic or with minimal human interaction on parameter selection. For example, in the surface segregation step (ii), a parameter d_{med} is automatically estimated by taking the median of inter-sampling distance from the full MS [CLK09]. In the regularisation step (iii), the stopping criteria is obtained by counting the number of sampling points n_s associated with each MS loop (or MA sheet) and selecting a threshold (on the maximum object surface size) either automatically by histogram analysis or imposed manually [CK08]. Other parameters are similarly either estimated from the data, or can be left set by the user.