



Fitting Unbranching Skeletal Structures to Objects

Zhiyuan Liu^{a,*}, Junpyo Hong^b, Jared Vicory^c, James N. Damon^d, Stephen M. Pizer^a

^aDepartment of Computer Science, University of North Carolina at Chapel Hill, USA

^bGE Healthcare, USA

^cKitware Inc., USA

^dDepartment of Mathematics, University of North Carolina at Chapel Hill, USA

ARTICLE INFO

Article history:

Keywords: Shape representation, Skeletal model, Statistical shape analysis

ABSTRACT

Representing an object by a skeletal structure can be powerful for statistical shape analysis if there is good correspondence of the representations within a population. Many anatomic objects have a genus-zero boundary and can be represented by a smooth unbranching skeletal structure that can be discretely approximated. We describe how to compute such a discrete skeletal structure (“d-s-rep”) for an individual 3D shape with the desired correspondence across cases. The method involves **fitting** a d-s-rep to an input representation of an object’s boundary. A good fit is taken to be one whose skeletally implied boundary well approximates the target surface in terms of low order geometric boundary properties: 1) positions, 2) tangent fields, 3) various curvatures.

Our method involves a two-stage framework that first, roughly yet consistently fits a skeletal structure to each object and second, refines the skeletal structure such that the shape of the implied boundary well approximates that of the object. The first stage uses a stratified diffeomorphism to produce topologically non-self-overlapping, smooth and unbranching skeletal structures for each object of a population. The second stage uses loss terms that measure geometric disagreement between the skeletally implied boundary and the target boundary and avoid self-overlaps in the boundary. By minimizing the total loss, we end up with a good d-s-rep for each individual shape. We demonstrate such d-s-reps for various human brain structures. The framework is accessible and extensible by clinical users, researchers and developers as an extension of SlicerSALT, which is based on 3D Slicer.

© 2021 Elsevier B. V. All rights reserved.

*Corresponding author: Tel.: +1-984-215-7219; fax: +1-919-590-6105;
e-mail: zhiy@cs.unc.edu (Zhiyuan Liu)

1. Introduction

3D anatomical shapes in medical image analysis are most intuitively understood in terms of their boundaries. However, analysis of objects' shapes benefits from models that capture shape features additional to boundary positions. For instance, objects' widths, boundary normals and boundary curvatures may be useful features (Srivastava and Klassen, 2016; Pizer *et al.*, 2019). The Blum medial representation (Blum, 1964; Siddiqi and Pizer, 2008) captures these features via a skeleton and its connection to points of spherical bitangency on the boundary. We call these connection vectors, *spokes*. As discussed in (Damon, 2008), this skeleton must share non-self-overlap topology with the object's boundary.

Medical images reflect aspects of structures but also include noise from imaging devices. Moreover, a manual or automatic discrete segmentation induces unintended deviations from true boundaries of objects. Together these pose notable challenges for anatomical shape analysis. Under a small perturbation of low order geometric properties, a skeletonization method (Kimia *et al.*, 1995; Siddiqi *et al.*, 1999) that derives a Blum medial structure from a boundary could result in a skeleton of branching topology that differs from that of the noise-free object, as illustrated on the left of [fig. 1](#). Such sensitivity to perturbations brings unnecessary variations and interferes with a succeeding statistical analysis.

To address the problem, Damon (2003) introduced the notion of a representation by radial flow of a skeletal model, in which the skeleton flows to the boundary along spokes, as distinguished from the previous models of boundary ("grassfire") flow to the skeleton, the behavior of which is mathematically studied in the work of Giblin and Kimia (2001). Based on radial flow, Damon introduced the notion of a general skeletal structure for an object that relaxes the appropriate Blum conditions while developing the mathematics needed for the structure to capture the geometric and topological properties of an object. Further, Pizer *et al.* (2019) proposed to use these relaxed conditions in fitting a smooth unbranching skeletal structure (an unbranching "s-rep") to an unbranching object boundary, as op-

posed to deriving the structure from the boundary (see [fig. 1](#)). Starting from a discretely sampled version of this s-rep ("d-s-rep"), the implied region generated by the radial flow (e.g., the region bounded by the dotted curve on the right of [fig. 1](#)) needs to have geometric features well matched to the target object, to be non-self-overlapping and to correspond well across a population of similar shapes. To achieve all three aims, we have devised a two-stage d-s-rep fitting framework: first, we initialize a d-s-rep that has consistent skeletal branching topology within the population; second, we refine the d-s-rep with respect to geometric match quality.

The first stage of the fitting process takes the topology of the population of objects into account. Many human structures are diffeomorphic to ellipsoids; namely, there exists a bijection that is differentiable with differentiable inverse and that maps the boundary of such structure to that of an ellipsoid. Moreover, many of them have no branching subfigures. When the ellipsoid is eccentric, we refer such object as a *slabular* object; roughly, its length is greater than its breadth, which is greater than its width. We choose the ellipsoid as the basic slabular object because 1) its Blum medial structure is an s-rep with a flat, unbranching structure whose edge is an ellipse that directly maps to the crest of the ellipsoid; 2) the map between its boundary and its skeletal representation has an analytic form; 3) it is easy to parameterize both the boundary and the skeletal sheet in such a way that the crest curve on the boundary of the ellipsoid divides it into two open halves, namely, a north and a south side; 4) there is a natural correspondence between eccentric ellipsoids via the ratios of their largest, middle-sized, and smallest principal radii respectively.

We infer the map between the target object and its skeletal representation given 1) a diffeomorphic map \mathcal{A} between the boundary of the target object and that of an ellipsoid and 2) the s-rep of the ellipsoid, from which a d-s-rep can be analytically computed. We extend \mathcal{A}^{-1} to the skeletal sheet, producing what is called a stratified diffeomorphism \mathcal{A}' . We apply \mathcal{A}' to the ellipsoid's d-s-rep to produce the initialized fit. We detail this stage in section 3.




Fig. 1: Left: a 2D object with simple Blum branching topology except for a boundary bump induced by noise. The noise induces an extra branch on the skeleton (blue). Right: a fitted skeleton to the object that obtains the noise-free object’s branching topology. This skeletal model is associated with a slightly modified (dotted) boundary. We aim to obtain such robust skeletal representations.

The second stage involves the improvement of the initial fit. The initial fit needs such improvement because it may result in a self-overlapping skeletal structure and may differ unnecessarily in geometric properties from the target objects. We measure the goodness of the fit in three low order geometric properties: the boundary positions, the tangent fields and the radial curvature (Damon, 2008) – respectively, zero-, first- and second-order geometric properties. Section 5 discusses the methodology of this refinement.

The refinement stage needs a finely sampled s-rep, while the initialization stage outputs a coarsely sampled d-s-rep. To fill this gap, section 4 describes the methodology of interpolating a d-s-rep into a more finely sampled d-s-rep or even a continuous s-rep.

This paper is organized as follows. In section 2 we review the related work. To demonstrate the proposed framework, we show results from experiments on an infant brain imaging study (IBIS) dataset. In section 6 we describe the data. Section 7 focuses on the experimental setup and the implementation details. Section 8 shows the results from both the initialization and refinement stages. Section 9 discusses the accomplishments of this paper and proposes future work. In order to make the framework publicly accessible, we have implemented it as an extension to 3D Slicer. It has been packaged in SlicerSALT (Vicory *et al.*, 2018) and was released recently.

2. Background

The skeletal fitting method described here differs from previous methods as surveyed (Siddiqi and Pizer, 2008). These

include a ridge following approach (“cores”), a boundary erosion approach, discrete morphology approaches, Voronoi approaches, continuous medial fitting approaches (“cm-reps”) (Yushkevich *et al.*, 2006; Pouch *et al.*, 2012) and our own laboratory’s previous fitting approach (Joshi *et al.*, 2002; Pizer *et al.*, 2003). Because our previous approach reflected some – but not all – of the Blum medial properties, we called the resulting representation an “m-rep”. None of these approaches, with the possible exception of a very recent form of cm-rep (Yushkevich *et al.*, 2019), can achieve the goal of automatically computing skeletal representations suitable for statistical analysis as discussed above.

Yushkevich *et al.* (2019) modified a method due to Arguillère *et al.* (2015) to construct a diffeomorphism between objects that mapped a finite set of points representing the medial axis and the corresponding boundary points for the first object to another finite set for the second object. Nonetheless, without taking into account C^1 or higher order information, such a diffeomorphism could fail to satisfy higher order conditions on the boundary as well as on the skeleton. As a result, this diffeomorphism can change the skeletal topology as shown in Damon (2021).

Hong (2019) showed the promise of automatically fitting a skeletal structure to a slabular anatomical structure; that is, a skeletal fitting method not needing a pre-defined skeletal template. Nonetheless, such a fit can result in a self-overlapping skeletal structure and can not strongly enough match the target object’s geometry (for example, see the middle and right panels of fig. 2). A refinement stage is needed.

Pizer *et al.* (2013) refined an m-rep within the shape space obtained from statistical analysis in order to achieve a tight shape distribution for a population. However, that shape space is free from geometric constraints, which still cannot resolve self-overlaps. Later, the refinement algorithm studied in Tu *et al.* (2016) takes into account both global consistency and individual specificity of m-reps. However, the algorithm relies on statistics (i.e., variation modes) of a population as the refinement proceeds. It is computationally expensive and impractical for refining each single medial structure.

Despite the above-mentioned drawbacks, these medial structures have been widely applied in anatomical shape analysis. For example, Tu *et al.* (2016, 2018) have demonstrated that entropy-based s-reps improve the performance of statistics, as compared to boundary models. Schulz *et al.* (2016) studied hypothesis tests of individual geometric object properties of s-reps. They found that using such geometric properties on hippocampi could effectively differentiate first episode schizophrenics and controls. Vicory *et al.* (2014) showed effective segmentations using s-reps. Also, Hong *et al.* (2016) have applied s-reps in classification tasks on two brain structures and on two diseases.

3. Initialize an s-rep

In this section, we first give a formal definition of an s-rep in both the continuous and discrete senses. Then we discuss the initialization of s-reps with consistent branching topology for slabular objects.

3.1. Skeletal representation

In this paper we restrict ourselves to s-reps of the form (M, S) , where M is formed from two smooth colocated sheets and thus has ellipsoidal topology, and where S is a field of non-crossing spoke vectors on M (see fig. 2 left). That is, (M, S) is formed from 3 components: (M_+, S_+) , (M_-, S_-) and (M_0, S_0) , where M_+ , M_- , M_0 compose the colocated skeletal sheets (in 3D, curves in 2D): M_0 is a fold curve of M that divides M into two open submanifolds, namely M_+ and M_- . M_+ is the south side of M , and M_- is the north side of M . Via S_+ , M_+ is mapped

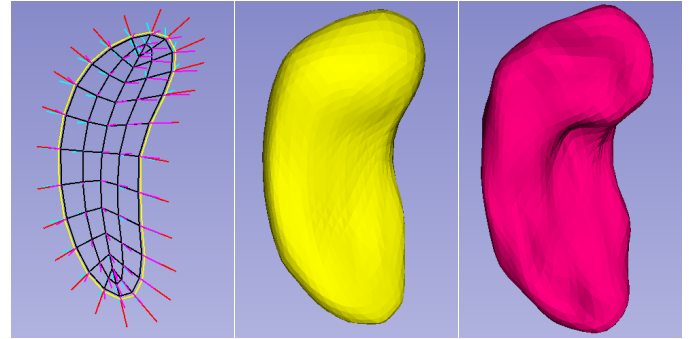


Fig. 2: Left: a 3D s-rep for a hippocampus. The grid connects skeletal points. The yellow curve is the skeletal fold. The vectors proceeding from the skeletal points are the spokes. Middle: the 3D implied boundary of this s-rep. Right: the surface mesh of this hippocampus being fitted.

to the southern boundary of the object; via S_- , M_- is mapped to its northern boundary; via S_0 , M_0 is mapped to a crest curve on its boundary.

S_+ , S_- and S_0 are continuous vector-valued functions; each vector $s \in S$ has a length r and direction U , and it proceeds from a point on M_+ , M_- or M_0 , respectively. The “spoke” s points from a skeletal point p to a point $\mathcal{B}(p)$ on the boundary.

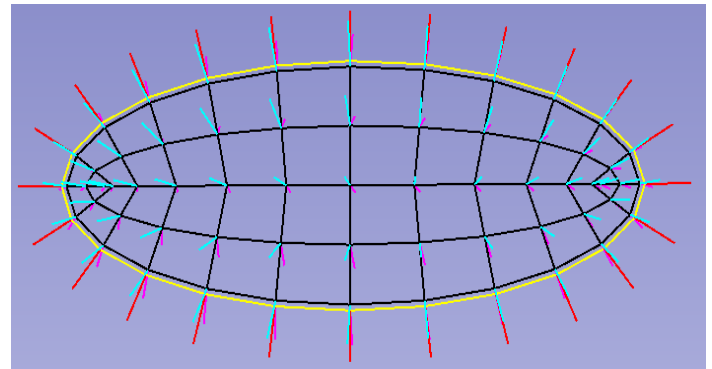


Fig. 3: An oblique view of a skeletal representation, where the skeletal sheet is discretized as a grid, of an ellipsoid. The interior black grid is the discretized colocated skeletal sheets M_l . Each grid point on M_l is associated with a spoke, pointing to either the north (magenta) or south (cyan) side of the ellipsoid's boundary. The yellow curve is the fold M_0 , from which red spokes point to the crest curve of the ellipsoid.

In Blum's view of medial models and Damon's view of skeletal models (Damon, 2003), the two colocated sheets M_+ and M_- are thought of as a single manifold M with boundary ∂M (i.e., M_0). The locations on $M \setminus M_0$ are the “smooth” points and the locations on M_0 are the fold (“edge”) points. In this conventional view, associated with each smooth point p there are two spokes $s_{+,p}$ and $s_{-,p}$. Notice that in a skeletal model the length of $s_{+,p}$ is not necessarily equal to that of $s_{-,p}$. Moreover, the

spokes swing infinitely fast as their base points approach and pass around the fold. We will detail this in section 5.

To establish correspondence among 3D objects, we discretize M_I (i.e., $M_+ \sqcup M_-$) into grids and M_0 into intervals, as shown in fig. 3. As a consequence, functions S_+ , S_- , S_0 are also sampled at each skeletal point. We denote the result from the discretization as a “d-s-rep” $(\widehat{M}, \widehat{S})$.

3.2. Automatic initialization methodology

Since many human organs have a similar boundary and skeletal topology with a slabular object, we can treat them as different derivations from a symmetric slabular object such as an ellipsoid (Leyton, 1992). Such uniform boundary and skeletal topology yields statistics-friendly (mutually corresponding) skeletal representations. Inspired by this, we establish an initial fit to a hippocampus by deforming a skeletal representation of an ellipsoid to that of the hippocampus.

To automatically obtain the corresponding ellipsoid and perform the deformation, we estimate stratified diffeomorphisms on boundaries via mean curvature flow¹ and on skeletons via the thin plate splines algorithm. Therefore, for each object we initialize a skeletal representation in three steps:

1. We establish the correspondence between the surface of the object and that of an ellipsoid (Hong, 2019).
2. We derive a continuous skeletal representation of the ellipsoid and discretize it.
3. We infer the interior mapping from the ellipsoid to the object and apply it on the discrete skeletal representation of the ellipsoid.

There is an easily computed smooth normal vector field on the boundary. To establish the correspondence between the object and an ellipsoid, we continuously perform mean curvature flow (Taubin, 1995) on the boundary at small steps, given the normal vector field. After each step the normal vector field is updated on the deformed boundary. We keep track of these local diffeomorphisms at each step. Consequently, the final map

from the object to the ellipsoid is the composition of these local diffeomorphisms.

While typically this mean curvature flow ultimately converges to a sphere, on the way it becomes a near-ellipsoid, denoted as \widetilde{E} . We measure the distance between the deformed boundary and the boundary of the best-fitting ellipsoid during the flow. We stop the flow once the distance is adequately small.

Next, we derive the skeletal representation for that best-fitting ellipsoid. Let r_1, r_2, r_3 be radii of the ellipsoid. Then any point (x, y, z) within or on that ellipsoid satisfies

$$\frac{x^2}{r_1^2} + \frac{y^2}{r_2^2} + \frac{z^2}{r_3^2} \leq 1 \quad (1)$$

where we assume $r_1 > r_2 > r_3$.

The flat skeletal (here medial) sheets M_I reside in the middle of the ellipsoid, and M_0 is an ellipse. Let a skeletal point $p \in M$ have coordinates (x_p, y_p, z_p) and let m_1, m_2 be the radii of M_0 .

Then we have

$$\frac{x_p^2}{m_1^2} + \frac{y_p^2}{m_2^2} \leq 1 \quad (2)$$

where

$$m_1 = \frac{r_1^2 - r_3^2}{r_1} \quad (3)$$

$$m_2 = \frac{r_2^2 - r_3^2}{r_2} \quad (4)$$

z_p is a constant. Let $z_p = 0$ for convenience.

We must now discretize the skeletal sheets of the ellipsoid. The location of the skeletal sample points was motivated by two considerations. First, the quality of the fit of the model into any object depends especially heavily on the fit at the crest and especially at its vertices (local extrema of convex curvature). For an ellipsoid these object surface loci correspond to the fold (end curve) of the medial ellipse and to that curve’s two vertices, respectively. Thus, as illustrated in fig. 3, on the medial ellipse we generate samples most densely at and near the fold and more densely along the fold near the vertices of the fold.

The second consideration is that, for consistency with the skeletal modeling of 3D objects, the sampling of the 2D medial ellipse should also be skeletal, or more precisely, medial. That is, the samples should be arranged along normals to the ellipse’s boundary, i.e., the fold curve M_0 . These normals are

¹Here, we assume the target boundary will not develop any singularities due to the flow as discussed in section 8. This assumption holds for structures such as the hippocampus, caudate, lateral ventricle, etc..

sampled at the boundary, at a position slightly eroded from the boundary, and then equally spaced between that position and the ellipse's medial axis (including a point on the medial axis). The rate of sampling along the normals is a parameter choosable by the user.

We sample the medial points along the fold curve (the yellow curve in fig. 3) as follows. The points on the fold are calculated as $(m_1 \cos\theta, m_2 \sin\theta, 0)$, and θ is uniformly sampled. The result is that the arclengths between adjacent fold points in the flatter region of the ellipse are notably greater than those near the vertex.

At each of these sampled points on the fold, we compute the interior sample points (the vertices of the black grid in fig. 3) as locations along the ellipse's normals at the respective fold points. The resulting grid is made up of quadrilaterals. Those near the ends of the ellipse's medial end-curve are quite small, and four of the quadrilaterals are degenerate, having two collinear sides.

We now must compute the spokes at each skeletal sample point. For each point $p \in M_I$, we find the nearest boundary point $\mathcal{B}(p)$ on the corresponding side of the ellipsoid analytically. p and $\mathcal{B}(p)$ form the endpoints of a spoke. Having both boundary points and corresponding skeletal points, we have formed the discrete skeletal representation $(\widehat{M}, \widehat{S})$.

Finally, we deform the discrete skeletal representation of the best-fitting ellipsoid to the target object. Because the boundary of the near-ellipsoid \widetilde{E} can be arbitrarily close to that of the best-fitting ellipsoid, this discrete skeletal representation can as well fit \widetilde{E} . Having the local diffeomorphisms in the process of flowing the boundary points of the object to those of \widetilde{E} , now we apply the inverse of these diffeomorphisms in a reverse order step by step on the boundary points of \widetilde{E} . Meanwhile, we infer the deformation on M via the thin plate splines algorithm and apply it to give the base points of spokes in \widehat{S} . This yields an initial fit of an s-rep to the target object.

4. Interpolation

This section discusses the methodology of interpolating a discrete skeletal representation (“d-s-rep”) into a continuous

s-rep in 3D. Given a d-s-rep, we treat the interior elements (i.e., skeletal points and spokes) and others (i.e., fold points and fold spokes) differently as follows.

4.1. Interpolate skeletal sheets and associated spokes

Each quadrilateral on a skeletal sheet is parameterized by (u, v) coordinates. We fit a cubic Hermite patch (Han et al., 2006) to each quadrilateral and interpolate it into a continuous patch. We interpolate in a *quad-by-quad* manner on M_I .

Then we interpolate both the directions and radii of the spokes via a successive subdivision method. The principle of interpolating directions is to have a smooth unit vector field on the respective open submanifold M_+ or M_- . Given a quadrilateral bounded by 4 skeletal points (each associated with a spoke), we interpolate unit vectors at centers of edges and the center of the quadrilateral by spherical linear interpolation (Shoemake, 1985). Along an edge starting from p_0 and ending with p_1 , we wish to interpolate between the orientation of the spoke at p_0 and that at p_1 . The associated orientation at λ (e.g., $\frac{1}{2}$) fraction away from p_0 to p_1 is given by the quaternion

$$q = q_0(q_0^{-1}q_1)^\lambda \quad (5)$$

where q_0, q_1 are quaternions associated with p_0 and p_1 , respectively. For the associated orientations at the center points in a two-dimensional quadrilateral, we interpolate the two dimensions individually in a similar fashion.

We apply this interpolation to produce sub-grid quadrilaterals of half size in each direction, producing a subdivision step. We repeat this process, yielding our successive subdivision for spoke directions.

We also use a similar successive subdivision strategy to interpolate the spoke radii, using skeletal geometry. Specifically, the associated radius at a middle point is given by

$$r_{p_{\frac{1}{2}}} = U_{p_{\frac{1}{2}}} \cdot \left(\frac{1}{2}(S_{p_0} + S_{p_1}) \right) - \frac{d^2}{4}(S_{p_0} \cdot U_{ww,p_0} + S_{p_1} \cdot U_{ww,p_1}) \quad (6)$$

where subscripts $p_0, p_{\frac{1}{2}}, p_1$ are respectively the start, middle (interpolated) and end positions. d is the distance from $p_{\frac{1}{2}}$ to p_0 . $U_{p_{\frac{1}{2}}}$ is the interpolated spoke direction at position $p_{\frac{1}{2}}$, and

$S_p = r_p \cdot U_p$. Further, U_{ww,p_0}, U_{ww,p_1} are the second derivatives of the spoke directions along the interpolating direction w at point p_0, p_1 , respectively. The first term can be understood as the average of two end spokes projecting to the spoke direction at the middle position. This projection should be a good estimate of the interpolated radius. The second term corrects this estimate based on the skeletal geometry. We refer the reader to [Pizer et al. \(2019\)](#) for a more rigorous explanation. For the center of a sub-grid quadrilateral, we interpolate the radii first along u then v and also along v then u , and we average the results.

4.2. Interpolate spokes in a crest region

For skeletal points on the fold curve, we interpolate interval by interval via the cubic Hermite splines.

The associated spokes' directions are differentiable along the fold curve ([Damon, 2008](#)). Thus, we can use the same interpolation strategy along the fold as described in section 4.1 and obtain the spokes associated with the fold.

Our definition of s-reps in section 3.1 allows spokes to swing smoothly as approaching and passing around the fold. In doing so, we fit splines on the boundary across the crest, and we interpolate along those splines to produce tips of interpolated spokes. Correspondingly, we interpolate the skeletal points between the outermost interior skeletal points on M_+ to corresponding fold points on M_0 and onto the outermost interior skeletal points on M_- . By doing so, we obtain interpolated spokes around the fold.

Were this whole interpolation to be applied continuously, it would turn a d-s-rep into a corresponding s-rep. Thus, assuming this interpolation, we can simply refer a d-s-rep as an s-rep. We will use the term ‘‘s-rep’’ in this way for the remainder of this paper.

5. Refinement

This section discusses the second stage of the fitting. While taking advantage of Damon's relaxation of medial conditions, we wish to deviate only slightly from the medial conditions while maintaining the non-self-overlap and branching topology. We accomplish this by designing an optimization in which we

penalize those deviations that can change most strongly in our initialization stage (see section 5.1). In sections 5.2-5.5 we address the penalties involving non-fold spokes. Then in section 5.6 we address the refinement on fold spokes. We defer the implementation details to section 7.

5.1. Relax Blum conditions

[Pizer et al. \(2013\)](#) have shown that the following enumerated objectives often provide a tighter fit to the geometry of the target object than a purely Blum-based fit. We will compare under two conditions the properties of the fit between the boundary \mathcal{B} of the target object and $\tilde{\mathcal{B}}$, that is created from the ends of the primary and interpolated spokes. The two conditions compared are when the full set of spokes satisfy the Blum conditions relative to $\tilde{\mathcal{B}}$ and when they are released from those conditions, according to Damon's skeletal geometry.

1. In both conditions \mathcal{B} and $\tilde{\mathcal{B}}$ do not precisely agree, even at the primary spokes. We do want the differences between \mathcal{B} and $\tilde{\mathcal{B}}$ to be small.
2. In both conditions the tangent planes of \mathcal{B} and $\tilde{\mathcal{B}}$ at corresponding spokes do not precisely agree, even at the primary spokes. We do want the differences between the tangent planes of \mathcal{B} and those of $\tilde{\mathcal{B}}$ to be small. However, in the case of Blum conditions the spokes are orthogonal to $\tilde{\mathcal{B}}$, i.e., they satisfy the partial Blum condition; Damon's skeletal geometry allows this condition not to hold.
3. According to ([Damon, 2003](#)), for a smooth point $p \in M_+$ and the collocated point $\bar{p} \in M_-$ in both conditions local self-overlap is prevented by the condition

$$r_p < \min\left\{\frac{1}{\kappa_{ri,p}}\right\}, \quad \text{if } \kappa_{ri,p} \in \mathbb{R}_*^+ \quad (7)$$

where r_p is the radius of the spoke associated with the point $\tilde{\mathcal{B}}(p)$ and $\kappa_{ri,p}$ is a positive real number denoting the i^{th} principal radial curvature at p . However, when the partial Blum condition holds, all of the eigenvalues of S_{rad} (the ‘‘radial shape operator’’, see section 5.4) are real, whereas in our relaxed skeletal situation, the condition need hold only for those places where the eigenvalues of S_{rad} are real. In our method, we wish this condition to

hold. We handle this in our objective function by a ReLU function (eq. (10)) described in section 5.4.

4. In both conditions a fold point p represents the maximum curvature point on the boundary at $\tilde{\mathcal{B}}(p)$. We wish $\frac{1}{r_p}$ to agree with the maximum curvature of the target boundary \mathcal{B} near $\tilde{\mathcal{B}}(p)$.

We also take advantage of another flexibility allowed by Damon’s skeletal geometry, namely, that the two spokes $s_{+,p}$ and $s_{-,p}$ need not be of the same length. However, because of the properties of our initialization stage, we have found it unnecessary to include a term as another objective making the difference in these spoke lengths small.

We approach the objectives 1–3 by penalizing deviations from those conditions with penalties L_0, L_1, L_2 . L_0 measures the overall squared distance from tips of spokes to the target boundary; L_1 measures the overall deviation of spokes’ directions from the perpendicularity to the boundary; L_2 measures the overall degree of violating objective 3.

It is the parameters of primary spokes $\widehat{\mathcal{S}}$, emanating from the grid points (i.e., \widehat{M}_I), that are refined. The interpolated spokes vary in response to this refinement. We evaluate the fit by measuring those penalties over the set of both primary and the corresponding interpolated spokes, which we call $\tilde{\mathcal{S}}$.

5.2. Penalize boundary distances

To compute the penalty L_0 , given the target surface mesh, we generate a 3D image of signed distances from the target object boundary. For each spoke $s_p \in \tilde{\mathcal{S}}$, we measure the distance d_{s_p} from its tip to the boundary according to the distance image. Thus, the total squared boundary distance penalty is

$$L_0 = \sum_{p \in \widehat{M}_I} \sum_{p' \in \mathcal{N}(p)} d_{s_{p'}}^2 \quad (8)$$

where \widehat{M}_I represents the sampled skeletal points and $\mathcal{N}(p)$ denotes the interpolated skeletal points near to the skeletal point p (see section 4).

5.3. Partial Blum condition

We penalize the non-orthogonality L_1 of spokes to a boundary. At each point of the signed distance image, we compute the

direction of the gradient which is perpendicular to the distance-level-set through that point. In particular, the gradient direction at a tip of a spoke is approximately perpendicular to the target boundary. We compare this gradient direction and the spoke direction, yielding the local penalty L_1 at p , i.e., $1 - \cos(\theta_p)$. Then we compute the overall penalty L_1 by

$$L_1 = \sum_{p \in \widehat{M}_I} \sum_{p' \in \mathcal{N}(p)} (1 - \cos(\theta_{p'})) \quad (9)$$

where $\theta_{p'}$ is the angle between the spoke direction $U_{p'}$ and the gradient direction at the tip of the spoke.

5.4. Resolve local self-overlaps

Objective 3 guarantees that we avoid local self-overlaps, and resulting singular kinks on the implied boundary, ensuring its smoothness. The principal radial curvatures $\kappa_{ri,p}$ are eigenvalues of the “radial shape operator” S_{rad} at $p \in M_I$ introduced by Damon (2003) for skeletal structures. This operator is a modification of the differential geometric shape operator (O’Neill, 2014), using instead the unit vector field in the spoke directions. This is usually not orthogonal to the medial sheet. Consequently, S_{rad} is not a symmetric operator, and if the partial Blum condition does not hold, its eigenvalues need not be real.

However, S_{rad} still captures the shape of the region bounded by the implied boundary; as well the regions bounded by any of the level surfaces of the radial flow (formed by dilating the skeletal sheet along each spoke by τ times its spoke length, for some $\tau \in [0, 1]$). These fill the region between the implied boundary and the skeletal set like a series of the onion-skins, which occur at specific times of the radial flow.

This operator S_{rad} has a matrix representation \widehat{S}_{rad} , which may not have real eigenvalues if the partial Blum condition doesn’t hold. We detail such a matrix representation for our s-reps in section 7.

For real principal radial curvatures at p from the eigenanalysis on \widehat{S}_{rad} , we can compute for each single spoke s_p of radius r_p the deviation from the objective 3 by a ReLU function:

$$\delta_{s_p} = \max(0, r_p \kappa_{r1} - 1) \quad (10)$$

where κ_{r1} is the maximal eigenvalue of \widehat{S}_{rad} at p . If instead the principal radial curvatures at p are not real, we let $\delta_{s_p} = 0$.

The overall L_2 penalty of an s-rep is

$$L_2 = \sum_{p \in \widehat{M}_1} \sum_{p' \in \mathcal{N}(p)} \delta_{s_{p'}} \quad (11)$$

5.5. Overall objective for interior spokes

The overall objective function for the refinement of interior spokes is given by

$$L = \alpha L_0 + \beta L_1 + \gamma L_2 \quad (12)$$

where weights α, β and γ are available to be set on a case-by-case basis. For example, one needs to tune up α and/or β if an initial fit was far from the object's boundary, and one needs to tune up γ if an initial fit presents self-overlaps. A good s-rep is obtained by minimizing eq. (12) over the parameters describing the primary spokes \widehat{S} .

5.6. Refine fold spokes

The initialization may result in fold spokes that violate objective 4; that is, fold spokes may deviate from the tangent plane of the s-rep. Our current solution to refine this condition is to adjust spokes' lengths, followed by relocating base points of fold spokes.

To do so, we start from adjusting lengths of fold spokes such that tips are on the boundary. Then, according to the maximal curvature $\kappa_{\mathcal{B}(p)}$ at a boundary point $\mathcal{B}(p)$, we relocate the skeletal end of this fold spoke so that the spoke's length becomes

$$r_p = \frac{1}{\kappa_{\mathcal{B}(p)}} \quad (13)$$

In shapes with sharp ends (e.g., the caudate nuclei), the relocation can make skeletal points of fold spokes far from the outermost interior ones (see fig. 4). This strategy can be viewed as extending M_I toward the sharp ends.

6. Materials

We collected MR brain images (the voxel size is $1mm \times 1mm \times 1mm$) from 177 6-month-old infants. From these images the subcortical structures were automatically segmented




Fig. 4: Fit an s-rep to a caudate nucleus in the dataset described in section 6. Top: the initial fit. Bottom: the refined s-rep. In the refined s-rep, the skeletal points of fold spokes (red line segments) are relocated according to curvatures of the crest.

(Wang *et al.*, 2014) to produce label images. We fit a surface triangular mesh to the boundary of each object of interest (i.e., a hippocampus or a caudate nucleus in this work) from each label image using SPHARM-PDM (Styner *et al.*, 2006).

7. Experimental settings and analysis

This section focuses on the settings and details of the experiments.

7.1. Initialize an s-rep for a hippocampus

Recall that the initialization stage has two stages. First, we flow the target shape boundary into a near-ellipsoid by following a series of local diffeomorphisms. We compute the best-fitting ellipsoid's s-rep that also fits to the near-ellipsoid. Second, we flow that s-rep back via a stratified diffeomorphism.

In the first stage, given the surface mesh of an object (see fig. 5 for an example of a hippocampus), we build a diffeomorphism from this boundary to that of a near-ellipsoid via mean curvature flow. At each step of the flow the boundary becomes smoother. To make the steps small enough and highly likely to be correct and easy to invert, it is recommended to perform 500 steps for a hippocampal shape. At the end the flow, we can achieve a near-ellipsoid (refer to fig. 5 middle) that is close




Fig. 5: Left: Input a 3D shape of a hippocampus represented by a surface mesh. Middle: The overlay of the target shape (red) and a near-ellipsoidal shape after mean curvature flow (black). Right: the best fitting ellipsoid (yellow transparent mesh) to the black mesh in the middle figure and its discrete s-rep (consists of cyan and magenta spokes, a discrete skeletal mesh and a fold curve).

enough to its best-fitting ellipsoid. We discretize the skeletal sheet of that best-fitting ellipsoid, sampling one point between the ellipse’s medial axis and the outermost interior skeletal point, shown on the right of fig. 5. This s-rep also well fits to the near-ellipsoid.

The second stage builds a stratified diffeomorphism to apply to the discrete skeletal points and to the corresponding spoke ends. This diffeomorphism is built from the vertices of the tile surface created on the target object. For the target objects (both the hippocampi and the caudate nuclei) in this paper, each surface mesh has 1002 vertices. With the stratified diffeomorphism we transform the s-rep of a near-ellipsoid to that of the target object, as shown in figs. 6 (a-c). Due to the stratified diffeomorphism, both the boundary match and the skeletal topology are preserved.

7.2. Refine the s-rep

We interpolate the discrete s-rep to evaluate the penalties L_0, L_1 and L_2 (refer to eqs. (8), (9) and (11)) as the primary spokes being refined in the procedure. An interpolated example is shown in fig. 6 (d).

We compute the loss term L_2 as follows. For each smooth point $p \in \widehat{M}_I$, we construct and compute the eigenvalues of the matrix $r\widehat{S}_{rad}$ (Han, 2008) representing the local radial shape at p as

$$r\widehat{S}_{rad} = \left(\left(\frac{\partial S}{\partial w} - \frac{\partial r}{\partial w} U \right) Q^T (Q Q^T)^{-1} \right)^T \quad (14)$$

where r and U are the length and direction of the spoke at p , respectively; w is a velocity direction of a curve on $T_p M$ through p , and

$$Q = \frac{\partial p}{\partial w} (U^T U - I) \quad (15)$$

where I is an identity matrix. As described in eq. (11), the eigenvalues of $r\widehat{S}_{rad}$ are used to form L_2 .

In fitting the hippocampi we set the weights α, β, γ in eq. (12) as 0.004, 20 and 50, respectively. It is recommended to subdivide each quadrilateral into a 8×8 sub-grid during the interpolation, yielding moderately interpolated s-reps. We found these settings also effective in fitting the caudate.

The optimization method that we use was chosen with the following considerations in mind:

1. The optimization can be over as few parameters as those describing a single primary spoke and as many as those describing the full set of primary spokes. So the method must be able to handle optimization over a few or over hundreds of parameters.
2. The objective function (eq. (12)) has many local optima, so a multiscale method is preferred.
3. No analytic expression for the objective function is available.
4. In each iteration of the optimization, as many as all of the primary spokes have been modified, so interpolation of many or all of the more finely sampled spokes must be

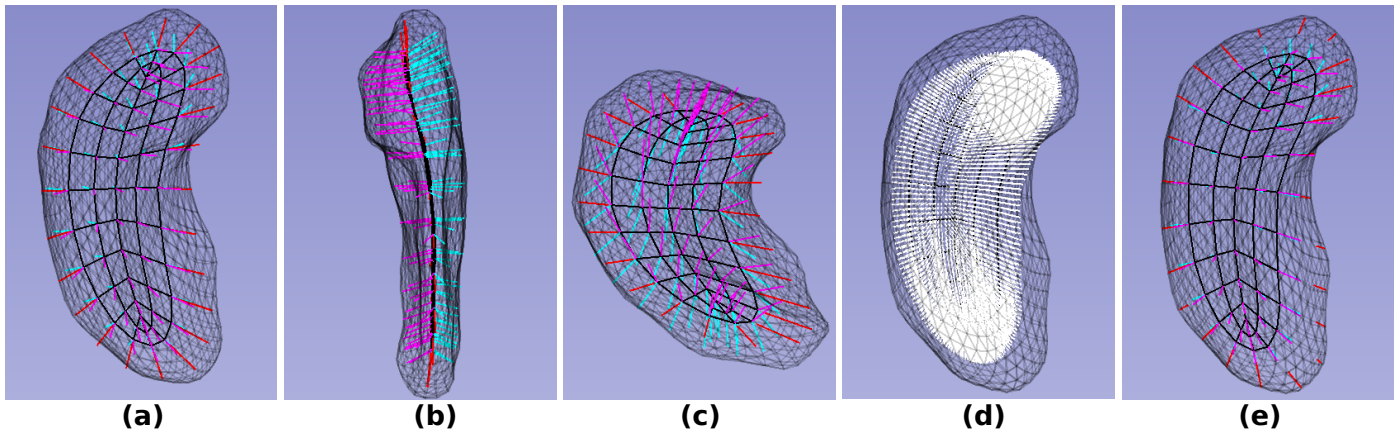


Fig. 6: (a-c) The result of the initial fit to a hippocampus from different perspectives. The gray transparent mesh is the input target surface mesh. The black grid in the middle of the object is a discrete representation of the skeletal sheet. The cyan and magenta spokes indicate spokes pointing to different sides. The red spokes are fold spokes. (d) The white spokes are the interpolated northern spokes of the initial fit. (e) The refined s-rep.

recomputed. Thus, computation of the gradient of the total loss by numeric approximation at every iteration would be unacceptably inefficient. Together, considerations 3 and 4 make a method involving gradients of the objective function inappropriate.

As a result, we have chosen a derivative-free optimization algorithm named “new unconstrained optimization algorithm” (NEWUOA) (Powell, 2004) to minimize the objective eq. (12).

8. Results and discussion

We tested our method on a machine with Ubuntu 18.04.3 LTS, 15.5 GB memory, Intel® Core™i7-8750H CPU @ 2.20GHz and GeForce GTX 1060. With the above-mentioned default parameters (e.g., the sampling rate of a skeletal sheet shown in fig. 3), a direct initialization takes about 2 minutes. A refinement from a previous result and the default parameters (e.g., the subdivision level) takes about 11 minutes.

We fit s-reps to the surface meshes of the hippocampi. Figure 7 shows the the initial and refined implied boundaries of a typical hippocampus. The initial implied boundary, shown in fig. 7(a), presents singular kinks in the middle, and these have been resolved by the refinement (see fig. 7(c)). The initial implied boundary has overall small distances from the target boundary positions (fig. 7(b)), while the refinement can reduce those distances around the crest region (see fig. 7(d)).

To examine the quality of the initial and refined fits, we compute the distances from tips of primary and interpolated spokes to the target boundaries at the recommended subdivision level. The statistics of these distances are shown in fig. 8. The results show that the worst case of the initialized s-reps has an object-average distance of $0.48mm$ between the implied and the target boundaries, as compared to $0.41mm$ after refinement. Due to the refinement, the object-average distance from implied boundaries to the target boundaries is improved by 11% on average. The improvement could be more significant if case-by-case parameters were employed. With slightly tuned parameters, we show the improvement of local distances between the implied and target boundaries in a typical example in figs. 8 (c) and (d).

With regard to the high order boundary geometric features, singularities on implied boundaries can be resolved by the refinement (see figs. 7 (a) and (c)). To quantify local shapes on the initial and refined implied boundaries as well as target boundaries, we measure the curvedness C (Koenderink, 1990)² on each boundary. The curvedness at a point on a surface is defined as

$$C = \sqrt{(\kappa_{max}^2 + \kappa_{min}^2)/2} \quad (16)$$

where κ_{max} and κ_{min} are the maximal and minimal curvature, respectively. C is close to 0 when the local shape is flat; C gets

²The corresponding notation used in the reference is R .

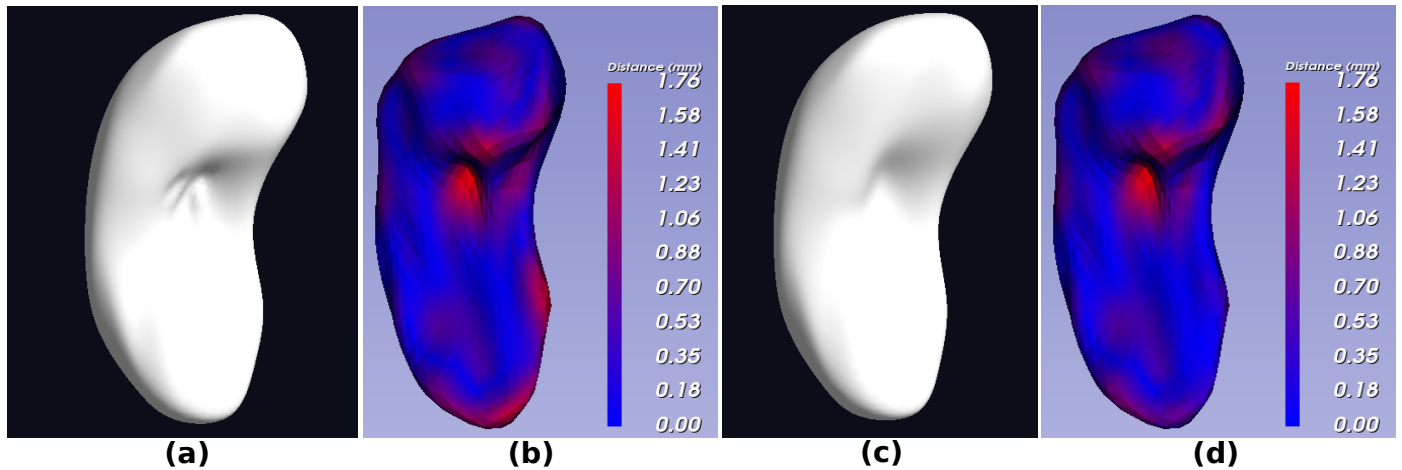


Fig. 7: (a): The implied boundary of the initial fit to a hippocampus. (b): The heat map of unsigned distances measured locally between the initial implied and target boundaries. (c): The implied boundary of the refined skeletal representation. (d): The local distances (in *mm*) from the target boundary to the implied boundary of the refined skeletal representation.

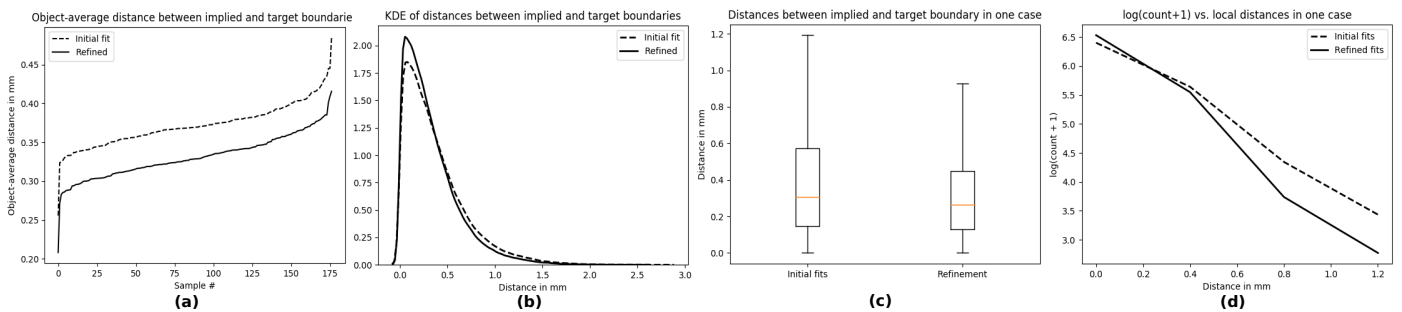


Fig. 8: (a): The object-average distances (measured in *mm*) between the implied and target boundaries in sorted order of 177 hippocampi. (b): The kernel density estimation (KDE) of histograms of average distances between the implied and target boundaries in these 177 hippocampi before and after the refinement. (c): the box plot of local distances from the target boundary in one typical case. (d): the histogram of the local distances in the same case.

larger when the local shape is sharp.

We present the histograms of local curvedness measured on an initial implied boundary, its refined implied boundary and its target boundary in [fig. 9](#). The smooth target boundary (see the dashed curve) resulting from SPHARM-PDM has a relatively tight distribution of curvedness. The initial fit (see the blue curve) implies many local patches with high curvedness on the boundary; the refinement smooths a lot of them (see the orange curve).

The pattern of points adopted to sample the skeletal surface was described in [section 3.2](#). The measurements used for [figs. 8](#) and [9](#) were made using this adopted pattern. Previously, we had made the hippocampal skeletal fittings and the same measurements using a pattern of points where the quadrilaterals were closer to being the same size ([Hong, 2019](#)), so as compared to

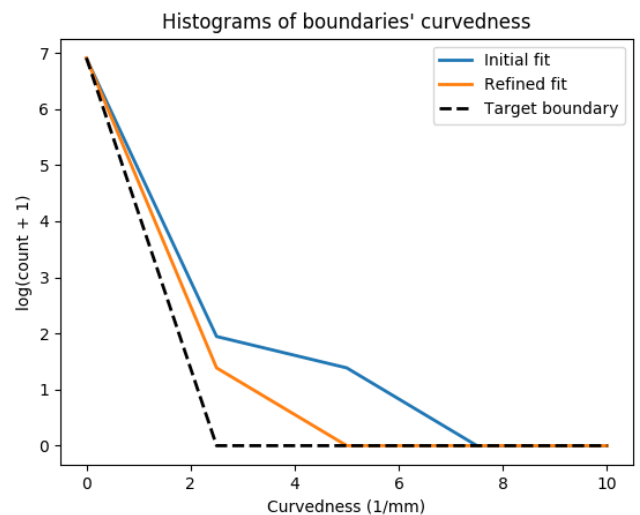


Fig. 9: The histogram of the curvedness on the initial and refined implied boundaries and target boundaries of a typical example. We apply the logarithm on the count of boundary point plus 1, yielding the y-value. The curvedness resulting from equation (16) is the x-value.

the adopted pattern, the central part of the basis skeletal ellipse was sampled more finely than with the final sampling pattern and it was sampled less finely near the vertices of the ellipse.

We have succeeded in fitting 3D s-reps to the caudate nuclei well (see [fig. 4](#) for example). We intend to apply the framework on other shapes.




Fig. 10: The conventional mean curvature flow develops singularities (e.g., around the circled region) for a mandible. Left: the input surface mesh of the mandible. Right: the result from a few steps of the flow.

However, the conventional mean curvature flow can develop singularities in evolving local shapes that have sections that are nearly rotationally symmetric and have convex mean curvature ([Huisken, 1990](#)), which happens, for example, with mandibles (see [fig. 10](#)). For such cases we would replace the flow with a modified mean curvature flow algorithm as studied in ([Kazhdan et al., 2012](#)), which is beyond the scope of this paper.

9. Conclusions and future work

The essential motivation of fitting s-reps is to obtain topologically consistent shape models that contain a rich collection of geometric features for statistical analysis. The fitting process should be automatic and independent on an explicitly pre-defined template. The implied boundary by a good s-rep should approximate the target boundary as to boundary positions, tangent fields and curvatures. To obtain such an s-rep, this paper has described a two-stage fitting framework: the first stage constructs a stratified diffeomorphism between an ellipsoid and the target object, resulting in an s-rep of desirable branching topology; the second stage refines the s-rep so as to minimize the deviations from the Blum conditions. As shown in the results, both initial fits and refined results approximate the boundary positions quite well. The refinement improves the fits of the

skeletal structures to the boundaries in terms of boundary distances and curvedness.

We have found that the fitted s-reps to the hippocampi and caudate nuclei have shown a good correspondence among objects. Also, the shape features from these s-reps have shown statistical significance in discriminating shapes of the autism against those of the non-autism group ([Pizer et al., 2019](#)).

We have packaged our algorithm in an open-source toolbox SlicerSALT (<http://salt.slicer.org/>). SlicerSALT provides the state-of-the-art statistical shape analysis methods as extensions of a customized 3D Slicer. It includes a number of powerful tools for visualizing and processing data. Also, the consistent user interfaces in SlicerSALT make it easy to operate. Like other extensions of 3D Slicer, SlicerSALT supports modules written in C++ and Python. A module in C++ is typically more efficient than that in Python. Thus, we implemented the initialization and the refinement modules using C++. In addition, we provided a visualization module in Python.

This method can also be applied on 2D skeletal modeling, i.e., to study the shapes of boundary contours in 2D images. To this end, we construct a generalized cylinder with smooth and closed boundary surface by repeating the target contour and capping it smoothly and symmetrically at both ends. Then we apply the framework discussed in this paper to fit a 3D s-rep to this generalized cylinder. An initial fit is obtained by slicing the 3D s-rep in the middle cross-section of the cylinder.

In future we will develop this method to fit s-reps to more complicated shapes (e.g., mandibles, patches of cerebral cortex, heart valves ([Pouch et al., 2015](#))). To do this, we will address the singularities developed in mean curvature flow. Also, we will continue investigating the strength of our s-reps in statistical shape analysis.

Acknowledgments

This work is funded by NIH R01EB021391. We thank Dr. Beatriz Paniagua at Kitware Inc. for her contribution in designing the initialization stage and her help and support during the implementation. We are also grateful for the support from Dr.

Martin Styner's group. We thank Dr. Joern Schulz, Mohsen Taheri and Jiyao Wang for their valuable feedback on the algorithm. We also have benefited significantly from Dr. J. S. Marron's advice during this work.

References

- Arguillère, S., Trlat, E., Trouv, A., Younes, L., 2015. Shape deformation analysis from the optimal control viewpoint. *Journal de Mathématiques Pures et Appliquées* 104, 139–178.
- Blum, H.F., 1964. A transformation for extracting new descriptors of shape. *Applications*. Springer. chapter Geometry and medial structure. pp. 69–123.
- Damon, J., 2003. Smoothness and geometry of boundaries associated to skeletal structures I: sufficient conditions for smoothness. *Annales de l'Institut Fourier* 53, 1941–1985.
- Damon, J., 2008. *Medial Representations: Mathematics, Algorithms and Applications*. Springer. chapter Geometry and medial structure. pp. 69–123.
- Damon, J., 2021. Rigidity properties of the blum medial axis. *Journal of Mathematical Imaging and Vision* 63, 120–129.
- Giblin, P., Kimia, B., 2001. On the local form and transitions of symmetry sets, medial axes, and shocks. *Proceedings of the IEEE International Conference on Computer Vision I*. doi:10.1023/A:1023761518825.
- Han, Q., 2008. Proper Shape Representation of Single Figure and Multi-Figure Anatomical Objects. Ph.D. thesis. Computer Science dissertation, Univ. of North Carolina at Chapel Hill.
- Han, Q., Pizer, S.M., Damon, J.N., 2006. Interpolation in discrete single figure medial objects, in: 2006 Conference on Computer Vision and Pattern Recognition Workshop (CVPRW'06), pp. 85–85.
- Hong, J., 2019. Classification of Neuroanatomical Structures Based on Non-Euclidean Geometric Object Properties. Ph.D. thesis. Computer Science dissertation, Univ. of North Carolina at Chapel Hill.
- Hong, J., Vicory, J., Schulz, J., Styner, M., Marron, J., Pizer, S., 2016. Non-euclidean classification of medically imaged objects via s-reps. *Medical Image Analysis* 31.
- Huisken, G., 1990. Asymptotic behavior for singularities of the mean curvature flow. *J. Differential Geom.* 31, 285–299.
- Joshi, S., Pizer, S., Fletcher, P.T., Yushkevich, P., Thall, A., Marron, J., 2002. Multiscale deformable model segmentation and statistical shape analysis using medial descriptions. *IEEE transactions on medical imaging* 21, 538–550.
- Kazhdan, M., Solomon, J., Ben-Chen, M., 2012. Can mean-curvature flow be modified to be non-singular?, in: *Computer Graphics Forum*, Wiley Online Library. pp. 1745–1754.
- Kimia, B.B., Tannenbaum, A.R., Zucker, S.W., 1995. Shapes, shocks, and deformations i: the components of two-dimensional shape and the reaction-diffusion space. *International journal of computer vision* 15, 189–224.
- Koenderink, J.J., 1990. *Solid Shape*. MIT Press, Cambridge, MA, USA. pp. 319–323.
- Leyton, M., 1992. *Symmetry, causality, mind*.
- O'Neill, B., 2014. *Elementary Differential Geometry*. Elsevier Science. URL: <https://books.google.com/books?id=HrriBQAAQBAJ>.
- Pizer, S.M., Fletcher, P.T., Joshi, S., Thall, A., Chen, J.Z., Fridman, Y., Fritsch, D.S., Gash, A.G., Glotzer, J.M., Jiroutek, M.R., et al., 2003. Deformable m-reps for 3d medical image segmentation. *International journal of computer vision* 55, 85–106.
- Pizer, S.M., Hong, J., Vicory, J., Liu, Z., Marron, J.S., et al., 2019. Object shape representation via skeletal models (s-reps) and statistical analysis. *Riemannian Geometric Statistics in Medical Image Analysis*, 233–272.
- Pizer, S.M., Jung, S., Goswami, D., Vicory, J., Zhao, X., Chaudhuri, R., Damon, J.N., Huckemann, S., Marron, J., 2013. Innovations for Shape Analysis: Models and Algorithms. chapter *Nested Sphere Statistics of Skeletal Models*. pp. 93–113.
- Pouch, A.M., A. Yushkevich, P., M. Jackson, B., S. Jassar, A., Vergnat, M., H. Gorman, J., C. Gorman, R., M. Sehgal, C., 2012. Development of a semi-automated method for mitral valve modeling with medial axis representation using 3d ultrasound. *Medical physics* 39, 933–950.
- Pouch, A.M., Tian, S., Takebe, M., Yuan, J., Gorman Jr, R., Cheung, A.T., Wang, H., Jackson, B.M., Gorman III, J.H., Gorman, R.C., et al., 2015. Medially constrained deformable modeling for segmentation of branching medial structures: Application to aortic valve segmentation and morphometry. *Medical image analysis* 26, 217–231.
- Powell, M.J.D., 2004. The newuoa software for unconstrained optimization without derivatives.
- Schulz, J., Pizer, S., Marron, J., Godtliebsen, F., 2016. Non-linear hypothesis testing of geometric object properties of shapes applied to hippocampi. *Journal of Mathematical Imaging and Vision* 54, 15–34.
- Shoemake, K., 1985. Animating rotation with quaternion curves. *SIGGRAPH Comput. Graph.* 19, 245–254.
- Siddiqi, K., Pizer, S., 2008. *Medial Representations: Mathematics, Algorithms and Applications*. 1st ed., Springer Publishing Company, Incorporated.
- Siddiqi, K., Shokoufandeh, A., Dickinson, S.J., Zucker, S.W., 1999. Shock graphs and shape matching. *International Journal of Computer Vision* 35, 13–32.
- Srivastava, A., Klassen, E.P., 2016. *Motivation for Function and Shape Analysis*. Springer New York, New York, NY. pp. 1–19.
- Styner, M., Oguz, I., Xu, S., Brechbhlher, C., Pantazis, D., Levitt, J., Shenton, M., Gerig, G., 2006. Statistical shape analysis of brain structures using spharm-pdm. *The insight journal* 1071, 242–250.
- Taubin, G., 1995. A signal processing approach to fair surface design, in: *Proceedings of the 22Nd Annual Conference on Computer Graphics and Interactive Techniques*, ACM, New York, NY, USA. pp. 351–358.
- Tu, L., Styner, M., Vicory, J., et al., 2018. Skeletal shape correspondence through entropy. *IEEE Transactions on Medical Imaging*.
- Tu, L., Vicory, J., Elhabian, S., Paniagua, B., Prieto, J.C., Damon, J.N., Whitaker, R., Styner, M., Pizer, S.M., 2016. Entropy-based correspondence improvement of interpolated skeletal models. *Computer Vision and Image Understanding*.
- Vicory, J., Foskey, M., Fenster, A., Ward, A., Pizer, S.M., 2014. Prostate segmentation from 3dus using regional texture classification and shape differences, in: *Proc of Shape 2014 - Symposium on Statistical Shape Models and Applications*.
- Vicory, J., Pascal, L., Hernandez, P., Fishbaugh, J., Prieto, J., Mostapha, M., Huang, C., Shah, H., Hong, J., Liu, Z., Michoud, L., Fillion-Robin, J., Gerig, G., Zhu, H., Pizer, S.M., Styner, M., Paniagua, B., 2018. Slicersalt: Shape analysis toolbox, in: *Shape in Medical Imaging - International Workshop, ShapeMI 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings*, pp. 65–72.
- Wang, J., Vachet, C., Rumble, A., Gouttard, S., Ouziel, C., Perrot, E., Du, G., Huang, X., Gerig, G., Styner, M.A., 2014. Multi-atlas segmentation of subcortical brain structures via the autoseg software pipeline. *Frontiers in neuroinformatics* 8, 7.
- Yushkevich, P.A., Aly, A.H., Wang, J., Xie, L., Gorman, R.C., Younes, L., Pouch, A.M., 2019. Diffeomorphic medial modeling, in: *Information Processing in Medical Imaging - 26th International Conference, IPMI 2019, Hong Kong, China, June 2-7, 2019, Proceedings*, pp. 208–220.
- Yushkevich, P.A., Zhang, H., Gee, J.C., 2006. Continuous medial representation for anatomical structures. *IEEE transactions on medical imaging* 25, 1547–1564.