

Continuous Medial Representations for Geometric Object Modeling in 2D and 3D

Paul Yushkevich, P. Thomas Fletcher, Sarang Joshi, Andrew Thall, Stephen M. Pizer

Medical Image Display and Analysis Group
University of North Carolina at Chapel Hill, USA
E-mail: pauly@cs.unc.edu

Abstract— We describe a novel continuous medial representation for describing object geometry and a deformable templates method for fitting the representation to images. Our representation simultaneously describes the boundary and medial loci of geometrical objects, always maintaining Blum’s symmetric axis transform (SAT) relationship. Cubic b-splines define the continuous medial locus and the associated thickness field, which in turn generate the object boundary. We present geometrical properties of the representation and derive a set of constraints on the b-spline parameters. The 2D representation encompasses branching medial loci; the 3D version can model objects with a single medial surface, and the extension to branching medial surfaces is a subject of ongoing research. We present preliminary results of segmenting 2D and 3D medical images. The representation is ultimately intended for use in statistical shape analysis.

I. INTRODUCTION

Medial loci, or skeletons, have enjoyed wide use in computer vision and medical image analysis because they provide important intuition about shape and formation of biological and anatomical objects. Medial loci naturally divide objects into a hierarchy of simple figures and describe the inherent symmetry and local thickness of each figure.

Medial loci of objects have traditionally been computed from discrete boundary-based descriptions by skeletonization algorithms. Such boundary descriptions, however, yield medial loci with a complex branching structure; for instance, each edge in a 3D polygonal boundary contributes a branch to the medial locus. Methods that simplify and regularize skeletons can eliminate unstable branches and yield object-relevant medial loci [18], [28], [23], [22]. Nevertheless, the boundary-to-medial transformation is inherently unstable; the resulting branching topology is sensitive to slight boundary perturbations, especially at the regions known as ligatures [4], [1].

Whereas the above methods start with a boundary description and yield the medial locus, *synthetic* medial representations, such as the one presented in this paper, use the medial loci themselves as a model for object representation. The model describes the medial branching topology and defines each branch of the medial locus using a few parameters. The model defines a smooth parameterized thickness field over the entire medial locus. A two-dimensional medial locus is a set of smooth curve segments joined at endpoints. A three-dimensional medial locus is a set of surface patches connected along curves. (See Fig. 1.)




Fig. 1. Branching medial loci in 2D and 3D. (a) In 2D, three branches generically join at shared endpoints. Approaching the shared endpoint, the thickness (radius) values associated with each branch converge to a common value. (b) In this 3D example, two medial surfaces are connected along a seam γ . The seam forms a crease in A and is part of the edge of B.

The medial locus and the associated thickness field synthesize a stable object boundary by inverting the skeletonization process. The generated boundary is equivalent to the envelope of spheres (or disks) placed at each point in the medial locus with the radius prescribed by the associated thickness value. The model establishes a correspondence between each point on the medial locus and a pair of points on the generated boundary. Synthetic medial representations enforce a fixed medial branching topology and provide a *simultaneous* description of the medial locus and the object boundary.

In this paper we present a continuous representation that uses cubic b-splines to model both the medial manifolds and the associated thickness field. We develop constraints on the parametric definition of the medial model that guarantee that the generated boundary surface is a closed, connected, non-singular, manifold with curvature continuity. The current state of the method allows representation of 2D objects with branching medial topology and 3D objects with a single medial surface. The extension to branching medial surfaces is the subject of ongoing research.

Our representation is essentially a continuous extension of *m-reps*, a discrete medial representation developed by Pizer et al. [26]. Discrete m-reps describe medial loci and boundaries of objects using sparse discrete samples called medial atoms. Each medial atom encapsulates local surface and thickness properties of the medial locus and implies local surface properties of the object boundary. M-reps have been shown effective for object representation, modeling, and image segmentation using deformable templates because natural operations such as bending, widen-

ing and elongation are easily implemented [17], [34], [36]. Styner automatically computed m-rep templates with fixed branching topology that through deformation accurately fit populations of hippocampi, amygdalae and other subcortical organs [32].

In this paper we show that the continuous medial representation can be applied to segment anatomical objects in medical images, following the same deformable templates framework used for discrete m-reps in work of Joshi, et al [17]. As an example, we automatically segment a vertebral image from a CT slice using a 2D model with branching medial topology. In 3D, we deform a template model of the hippocampus to fit manually segmented magnetic resonance images of the brain.

We developed continuous m-reps with the ultimate goal of improving our present methods in statistical shape analysis. The methods previously developed in our lab use discrete m-reps with a fixed branching topology to describe a population of objects [31], [39]. These methods estimate probability distributions of the statistical features derived from medial atoms. These distributions are used to generate new instances of m-reps, to visualize the primary modes of variability in the population in terms of bending or thickening of objects, to perform classification on the basis of shape and to pinpoint locations in objects where shape variability is most pronounced. Continuous m-reps augment shape analysis methods by allowing arbitrary sampling of medial loci. The continuous medial representation makes it possible to elastically model and optimize correspondences between features of different objects in the population.

The paper is organized as follows. In section II, related research in boundary and medial object representation is described. Section III develops the geometric foundation for continuous medial modeling and defines the generative medial b-spline model. Section IV describes the procedure that computes the optimal parameters of the model. Section V presents the preliminary results of fitting the model to images. Section VI describes the present limitations of the method, and compares the b-spline medial representation with discrete m-reps.

II. PRIOR WORK

This section summarizes the extensive prior work in areas of shape modeling and medial geometry. We loosely group the references into subjects of local boundary representations, global boundary representations, medial representations, and geometry of medial loci.

Local boundary representations, polygonal or higher order approximations of the boundary surface, are the principal shape descriptors in computer graphics, image analysis and shape characterization. These representations by themselves provide no *a priori* global information, such as the figural hierarchy of an object, and thus do not provide an intuitive framework for shape analysis and figure-based deformation. Nevertheless local primitives are used widely, for instance for elastic deformable modeling of 3D objects of arbitrary topology [10].

Global boundary shape descriptions have found use in

image analysis via deformable models and in shape analysis. Staib and Duncan represent boundaries of three-dimensional objects as weighted sums of fourier components [29]. Székely et al. use a spherical harmonics decomposition to represent three-dimensional objects of spherical topology [33]. Carr et al. use radial basis functions to represent boundaries [5]. Styner et al. combines the spherical harmonic SPHARM representation with discrete m-reps and computes statistics on the combined representation, yielding a comprehensive statistical analysis of shape [31].

Object description via skeletons, i.e. medial representation, has been used increasingly in computer graphics and CAD. Bloomenthal and Sherstuck develop implicit surfaces based on convolution over medial skeletons [3], [27]. Markosian et al. apply *skinning* of implicit fields around polyhedral skeletons [19]. Gascuel et al. develop a system for animation and collision detection based on rigid articulated skeletons fleshed by spline-based deformable boundary surfaces [12]. Storti et al. and Blanding et al. use a skeleton-based object representation for CAD style applications: 3D geometric model synthesis, generation of boundary surfaces at varying levels of detail, and morphing [30], [2]. Igarashi's Teddy system uses the medial spines to drive intuitive shape modeling based on hand sketching [16]. Chen uses multiscale medial models based on a sampled skeleton for guiding volume rendering [7]. Thall et al. propose discrete m-reps as a general geometric modeling primitive for 3D design in computer graphics and CAD, using displaced subdivision surfaces for fine-scale boundary description [36].

Use of medial representations for deformable modeling in computer vision was pioneered by work of Tsao and Fu, where discrete skeletons computed by distance transform are stochastically manipulated and a discrete boundary is regenerated [37]. Zhu and Yuille developed the comprehensive FORMS system that automatically divides 2D objects into simple parts and represents these parts medially, incorporating statistical shape information [40]. M-reps, a multiscale medial description on which this paper is based, and their use in medical image analysis have been discussed in section I. Our method differs from the ones above because we enforce the SAT relationship between a continuous deformable medial model and the boundary generated by it.

Differential geometric properties of medial axis transforms, boundary generation from continuous skeletons, and associated validity constraints have been studied extensively in the literature. In particular, Nackman [21], Vermeer [38], and Gelston [13] explored relationships between the 3D SAT and the curvature of its implied boundary. Hoffman & Vermeer [15], and Teixeira [35] described validity conditions that continuous medial surfaces must satisfy in order to exist and to imply non-intersecting boundaries. Giblin and Kimia studied local differential geometry of various kinds of points forming 3D symmetry sets [14]. Damon studied the geometry of offset surfaces generated by a multi-valued vector field defined on a set of connected manifolds; the results presented in the following section

represent a special case of Damon’s theory [8].

III. METHOD

This section presents the details of the medial b-spline representation. Section III-A describes the differential geometry of continuous medial manifolds. Section III-B presents the constraints that must be satisfied by a continuous medial model in order for the resulting boundary to form a closed non-singular manifold. In section III-C we implement the continuous medial representation, subject to the above constraints, using cubic b-splines.

The geometric notions presented in the first two subsections constitute a special case of the results recently developed by Damon. In [8] Damon studies a class of manifolds formed by offsetting a branching medial locus by a multi-valued vector field. Damon describes the differential geometry of the offset manifold in terms of the shape operator of the vector field and gives a set of necessary and sufficient constraints that both the medial surface and the vector field must satisfy in order to produce a continuous boundary without singularities.

A. Medial Geometry

In this work we describe a continuous medial representation based on the medial loci described by Blum as the symmetric axis transform (SAT)[4]. The SAT is constructed as a locus of centers of maximal disks inscribed into a geometric object. The local thickness of the object is described by the radii of the disks.

A synthetic continuous medial representation (*cm-rep*) defines the medial locus of an object as a set of connected parameterized manifolds called *medial manifolds*. A thickness value is associated with each point of each medial manifold.

Medial manifolds in 2D are smooth curve segments connected to each other at endpoints. Medial curve segments connect when three of them come together at a shared endpoint, as demonstrated in Fig. 1a. The thickness value is equal at the shared endpoint for all three branches. A joining of more than three branches at a point is non-generic.

In general, a 3D medial locus is a set of connected medial surfaces and space curves. Tubular objects whose medial loci are curves are non-generic and are not dealt with in this paper. Each medial surface is bounded by a closed curve that we call the *edge*. In particular, the part of an edge that is not part of any other medial surface is called the *free edge*.

Medial surfaces connect along shared curves that we call *seams*. Seams either form a part of the edge of a medial surface or form a crease (a discontinuity in the surface normal) on the medial surface. Medial surfaces are smooth except at seams. Fig. 1b, shows the seam γ that is a part of the edge of A and a crease on B . Geometric aspects pertaining to 3D cm-reps with connected medial surfaces are discussed below but the spline-based implementation is limited to models with a single smooth medial surface.

Formally, let \mathcal{O} be some geometric object with a closed continuous boundary. The cm-rep of \mathcal{O} is an approxima-

Symbol	Description
\mathbf{m}	Medial surface
r	Radial scalar field
u, v	Parametrization of (\mathbf{m}, r)
\mathbf{b}^t	Boundary counterparts of (\mathbf{m}, r)
t	Indexes the two parts $(-1, 1)$ of the implied boundary.
\mathbf{U}^t	Unit normal to the boundary, also the direction from a point on \mathbf{m} to its boundary counterpart.
\mathbf{N}	Unit normal to the medial surface.

TABLE I
MEDIAL NOTATION.




Fig. 2. Elements of 3D medial geometry. (a) The medial surface \mathbf{m} and its tangent plane at a point; the non-orthogonal frame $(\mathbf{m}_u, \mathbf{m}_v, \mathbf{N})$, and vector ∇r , the gradient of the thickness scalar field. (b) Implied boundary surfaces $\mathbf{b}^1, \mathbf{b}^{-1}$ and the vectors $\mathbf{U}^1, \mathbf{U}^{-1}$ that point from the medial surface to the boundary and are normal to the boundary.

tion of its medial axis. The cm-rep of \mathcal{O} is considered *valid* if it forms an exact medial axis of a geometric object that has the same topology as \mathcal{O} , and whose boundary is closed, continuous in curvature, and non-singular. This boundary is called the *implied boundary* of a valid cm-rep. The accuracy with which a cm-rep describes \mathcal{O} is measured in terms of differences between \mathcal{O} and the object formed by the implied boundary of the cm-rep.

In n -dimensional space, the medial locus is described by a set of control parameters that define C^2 functions $(\mathbf{m}, r) : D \rightarrow \mathbb{R}^n \times \mathbb{R}^+$ on a closed domain $D \subset \mathbb{R}^{n-1}$. The medial manifold is defined by the spacial component \mathbf{m} , and r defines the thickness field on \mathbf{m} .

The boundary generated by a cm-rep is constructed by inverting the SAT. Spheres (or disks) of radius $r(u, v)$ are placed at each location $\mathbf{m}(u, v)$ on the medial manifold. The generated boundary is the envelope of such a family of spheres or disks. In 3D, the points $\mathbf{x} \in \mathbb{R}^3$ that belong to this two-parameter family of spheres are defined by the implicit equation

$$f(\mathbf{x}, u, v) = |\mathbf{x} - \mathbf{m}(u, v)|^2 - r(u, v)^2 = 0. \quad (1)$$

At the points on the envelope f must satisfy

$$f = 0, \quad f_u = 0, \quad f_v = 0. \quad (2)$$

This system of equations yields a definition of the boundary that assigns a pair of boundary positions to each (u, v) .

Each sphere in the family generically touches the envelope at two points on opposite sides of the medial axis (with the exception special points where the two sides of the boundary come together, which will be discussed later). The points of tangency are called the *boundary counterparts* of a point on the medial axis. The two boundary counterparts are indexed by $t \in \{-1, 1\}$, are denoted as $\mathbf{b}^t(u, v)$, and are expressed as:

$$\mathbf{b}^t = \mathbf{m} + r\mathbf{U}^t, \quad (3)$$

$$\mathbf{U}^t = -\nabla r + t\sqrt{1 - \|\nabla r\|^2}\mathbf{N}, \quad (4)$$

where \mathbf{N} is the unit normal on the medial surface and \mathbf{U}^t is the unit normal on the boundary surface. The vector ∇r is the gradient of the thickness scalar field on the medial surface:

$$\nabla r = \begin{bmatrix} \mathbf{m}_u & \mathbf{m}_v \end{bmatrix} \mathbf{I}_m^{-1} \begin{bmatrix} r_u \\ r_v \end{bmatrix}, \quad (5)$$

where \mathbf{I}_m is the metric tensor on the medial surface.

The projections of both boundary counterparts onto the medial tangent plane lie in the negative ∇r direction. The distance from each counterpart to the medial tangent plane is $r\sqrt{1 - \|\nabla r\|^2}$; hence the boundary counterparts of \mathbf{m} are defined only if

$$\|\nabla r\| \leq 1. \quad (6)$$

This becomes the first constraint on the radial field.

The vector \mathbf{U}^t in the direction from a medial point to its boundary counterpart is normal to the implied boundary. This ensures that the medial surface is the SAT of its implied boundary. This property is used to express the curvature tensor of the implied boundary in terms of second derivatives of (\mathbf{m}, r) .

The second derivative tensor of the implied boundary surface is

$$\begin{aligned} \mathbf{II}_{\mathbf{b}^t} &= \begin{bmatrix} \mathbf{b}_{uu}^t \cdot \mathbf{U}^t & \mathbf{b}_{uv}^t \cdot \mathbf{U}^t \\ \mathbf{b}_{vu}^t \cdot \mathbf{U}^t & \mathbf{b}_{vv}^t \cdot \mathbf{U}^t \end{bmatrix} \\ &= - \begin{bmatrix} \mathbf{b}_u^t \cdot \mathbf{U}_u^t & \mathbf{b}_v^t \cdot \mathbf{U}_u^t \\ \mathbf{b}_u^t \cdot \mathbf{U}_v^t & \mathbf{b}_v^t \cdot \mathbf{U}_v^t \end{bmatrix}. \end{aligned} \quad (7)$$

The metric tensor of the implied boundary is

$$\mathbf{I}_{\mathbf{b}^t} = \begin{bmatrix} \mathbf{b}_u^t \cdot \mathbf{b}_u^t & \mathbf{b}_u^t \cdot \mathbf{b}_v^t \\ \mathbf{b}_v^t \cdot \mathbf{b}_u^t & \mathbf{b}_v^t \cdot \mathbf{b}_v^t \end{bmatrix}. \quad (8)$$

The principal curvatures and principal directions of the implied boundary, which are the eigenvalues and eigenvectors of $\mathbf{II}_{\mathbf{b}^t}\mathbf{I}_{\mathbf{b}^t}^{-1}$, are expressed in terms of first derivatives of \mathbf{b}^t and \mathbf{U}^t , and following (3), in terms of second derivatives of \mathbf{m} and r .

At a point \mathbf{m} on the medial surface where $\|\nabla r\| = 1$ the component of \mathbf{U}^t in the direction normal to \mathbf{m} is 0, following (4). The two boundary counterparts collapse to a common point on the tangent plane of the medial surface. The square root in (4) forces the derivative of \mathbf{U}^t to asymptote at such points, and the curvature tensor on the surface can not be computed in terms of second derivatives of \mathbf{m} and r .




Fig. 3. Examples of 3D cm-reps. (a) A medial surface. (b) Same medial surface and its implied boundary half \mathbf{b}^1 . (c) Same medial surface with both boundary halves $\mathbf{b}^1, \mathbf{b}^{-1}$ that form a closed surface.

The definition of the implied boundary in 2D is analogous to 3D. Equations (3), (4) do not change, and the vector ∇r is given by

$$\nabla r = \frac{dr}{ds}\mathbf{T} = \frac{r'}{\sqrt{x'^2 + y'^2}}\mathbf{T}, \quad (9)$$

where s is the arclength of the medial curve, \mathbf{T} is the unit tangent vector on the medial curve, and the primes are derivatives taken with respect to the parameter u .

B. Constraints on Medial Manifolds

In this section we describe two categories of constraints that the functions (\mathbf{m}, r) must satisfy. The first category ensures that the cm-rep generates a closed connected boundary. The second category ensures that the boundary is non-singular. We concentrate on 3D constraints and later point out how they are simplified in the 2D case.

Each medial surface comprising a cm-rep generates two boundary halves, \mathbf{b}^1 and \mathbf{b}^{-1} . We formulate local constraints that force the boundary halves to meet, forming a closed surface. These constraints differ at free edges, smooth interior points, seams and at the ends of seams. The spheres placed at these different classes of points have distinct orders of contact with the implied boundary, and have different geometric properties [14].

At free edges the boundary halves implied by the same medial surface connect. Recall from section III-A that the boundary counterparts of a point on a medial surface coincide when $\|\nabla r\| = 1$. Thus this condition must hold at points on the free edge. Fig. 4a shows an example of a medial surface that violates the free edge constraint.

At smooth points interior to a medial surface the two boundary counterparts are disjoint; otherwise the implied boundary does not form a closed manifold. Following (6), the constraint at these points is the strict inequality $\|\nabla r\| < 1$.

At a seam point three smooth subsurfaces $\mathbf{m}_i, i \in \{1, 2, 3\}$ come together at an angle to each other. For example, in Fig. 1b the opposite sides of the crease in A and the surface B form the three smooth subsurfaces. The six boundary halves implied by the three subsurfaces connect in such a way that \mathbf{b}_i^1 meets $\mathbf{b}_{i \oplus 1}^{-1}$ smoothly. The addition operation \oplus on the index i is cyclic on the set $\{1, 2, 3\}$. By




Fig. 4. Examples of boundary illegalities. (a) An example of a medial surface with an implied boundary that is not closed. (b) A swallowtail singularity on the implied boundary.

(3), the six boundary counterparts meet if

$$\mathbf{U}_i^1 = \mathbf{U}_{i\oplus 1}^{-1}, \quad (10)$$

since the r value is the same for all three subsurfaces at a seam point. Solving (4), (10) for \mathbf{N}_i yields the following constraint on seam points:

$$\mathbf{N}_i = \frac{\nabla r_{i\oplus 2} - \nabla r_{i\oplus 1}}{\sqrt{1 - \|\nabla r_i\|^2}}. \quad (11)$$

At seam endpoints one of the angles between the three joining subsurfaces becomes π and the constraint (11) disappears. Applicable free edge constraints must still be satisfied.

The above constraints ensure that the implied boundary is a closed and connected surface but do not guarantee that it is non-singular. Fig. 4b shows an implied boundary forming a swallowtail singularity, which commonly occurs when the medial surface is left unconstrained. Singularities and invalid regions are detected using the Jacobian of the mapping in (3), given by

$$J^t = t \frac{\mathbf{b}_u^t \times \mathbf{b}_v^t \cdot \mathbf{U}^t}{\mathbf{m}_u \times \mathbf{m}_v \cdot \mathbf{N}}, \text{ for } t = -1, 1. \quad (12)$$

At the singularities $J^t = 0$ and there exists a region where $J^t < 0$. We ensure that $J^t > 0$ at each point on the medial surface, thus eliminating singularities and points with reverse boundary orientation.

When the generated boundary is convex or hyperbolic at a point, the inscribed sphere must have smaller radius than the radii of curvature corresponding to the negative principal curvatures. Otherwise, the inscribed sphere would cross the implied boundary. Hence the constraint

$$r < -\frac{1}{\kappa_{min}^t} \quad (13)$$

must be satisfied by a valid cm-rep.

In two dimensions the above set of constraints reduces to a set of simple requirements. At free endpoints $|dr/ds| = 1$, while $|dr/ds| < 1$ is enforced on the interior of medial curves. At shared endpoints where three medial curves meet at an angle to each other, the following constraint holds:

$$\frac{dr_i}{ds_i} = \mathbf{T}_{i\oplus 1} \cdot \mathbf{T}_{i\oplus 2}, \quad (14)$$




Fig. 5. A manually constructed 3D b-spline medial of the hippocampus. (a) The medial surface and the control point polygon. (b) The implied boundary.

where the three tangent vectors are taken pointing into the shared endpoint. The Jacobian constraint from (12) simplifies to

$$\mathbf{b}_u^t \cdot \mathbf{T} > 0, \quad (15)$$

and the radius of curvature constraint from (13) does not change.

C. Generative Model

Cubic b-splines are used to model the medial surface and the thickness scalar field because they provide local control and C^2 continuity. The formulation of a smooth medial manifold as a b-spline is similar in 2D and 3D. We begin with the 3D case, which is presently limited to a single smooth medial surface.

A smooth medial surface $(\mathbf{m}(u, v), r(u, v))$ is defined in terms of control points as follows:

$$\begin{aligned} \mathbf{m}(u, v) &= \sum_{i=0}^{d_1} \sum_{j=0}^{d_2} N_i^3(u) N_j^3(v) \bar{\mathbf{m}}_{ij}, \\ r(u, v) &= \sum_{i=0}^{d_1} \sum_{j=0}^{d_2} N_i^3(u) N_j^3(v) \bar{r}_{ij} \end{aligned} \quad (16)$$

where $(\bar{\mathbf{m}}_{ij}, \bar{r}_{ij}) \in \mathbb{R}^3 \times \mathbb{R}$ is a $(d_1 + 1)$ by $(d_2 + 1)$ array of control points that include both positional and thickness components. N_i are third order b-spline basis functions [25].

Medial b-splines must satisfy the legality constraints defined in section III-B. The constraints ensuring that the generated boundary is closed are enforced differently in 2D and 3D.

B-spline surfaces are defined on a quadrilateral mesh and thus have sharp corners. It is not practical (or even possible) to enforce the edge condition $\|\nabla r\| = 1$ on the edge of such a rectangular surface, as it would limit the range of objects that could be represented.

Rather than enforcing the edge constraint along some predetermined curve on the b-spline surface, we restrict the medial surface to a subset that satisfies $\|\nabla r\| \leq 1$. The level curve $\|\nabla r\| = 1$ defines the edge of the medial surface. A b-spline surface must contain a single such level curve in its entirety in order to generate a boundary with spherical topology. By setting the \bar{r}_{ij} to large negative values at the perimeter control points, and by keeping \bar{r}_{ij} positive at the interior, we ensure that a level curve $\|\nabla r\| = 1$ lies on the b-spline surface. The large difference in \bar{r}_{ij} between




Fig. 6. Simple constraints on neighboring control points are used to ensure that the implied boundary is closed. (a) At free endpoints, the radius at the next-to-last control point is constrained. (b) At shared endpoints the control point $\bar{\mathbf{m}}_0$ is also shared and the radius value at the three neighboring control points is constrained.

adjacent control points causes the values of r_u and r_v to become large, in turn causing ∇r to become large, as a consequence of (5).

The level curve that forms the edge of a medial branch is computed by a numerical root finder in the process of sampling the spline. Whenever two consecutive samples have opposite signs of $\|\nabla r\| - 1$, Newton's method is applied along the vector connecting the two samples in parameter space. Fig. 3a is an example of a b-spline surface trimmed by the edge curve.

B-splines are an especially elegant representation for 2D cm-reps. Not only is it possible to represent medial structures with multiple connected branches, but also to incorporate all of the constraints on ∇r directly into the b-spline model. These constraints can be expressed as simple relationships between adjacent control points.

The knot sequences used to construct 2D b-splines have four repeating zero values at the beginning and four repeating unity values at the end. As a result, the b-spline behaves like a Bezier curve at the end points: it interpolates the terminal control points and the first derivatives of (\mathbf{m}, r) at curve ends are linear combinations of the terminal and the next-to-terminal control points, and shown in Fig. 6a.

At free endpoints, the constraint $|dr/ds| = 1$ is expressed in terms of the terminal control point $(\bar{\mathbf{m}}_0, \bar{r}_0)$ and the next control point $(\bar{\mathbf{m}}_1, \bar{r}_1)$:

$$\bar{r}_1 = \bar{r}_0 + \|\bar{\mathbf{m}}_0 - \bar{\mathbf{m}}_1\|. \quad (17)$$

At the points where three branches meet, the terminal control point of each branch is shared, as seen in Fig. 6b. The connectivity constraints from (14) are expressed in terms of the shared control point $(\bar{\mathbf{m}}_0, \bar{r}_0)$ and the next-to-last control points $(\bar{\mathbf{m}}_i, \bar{r}_i)$, $i \in \{1, 2, 3\}$ of each branch:

$$\bar{r}_i = \bar{r}_0 + \|\bar{\mathbf{m}}_0 - \bar{\mathbf{m}}_i\| \frac{(\bar{\mathbf{m}}_0 - \bar{\mathbf{m}}_{i\oplus 1}) \cdot (\bar{\mathbf{m}}_0 - \bar{\mathbf{m}}_{i\oplus 2})}{\|\bar{\mathbf{m}}_0 - \bar{\mathbf{m}}_{i\oplus 1}\| \|\bar{\mathbf{m}}_0 - \bar{\mathbf{m}}_{i\oplus 2}\|}. \quad (18)$$

Hence, a closed and connected generated boundary is formed by expressing the next-to-terminal control values \bar{r}_i in terms of near control points.

IV. PARAMETER ESTIMATION FOR IMAGE SEGMENTATION

Continuous medial models are used in image analysis applications following the deformable models framework.

A template cm-rep is deformed to optimally fit the image data in the presence of a geometric prior term.

Template models are constructed manually by moving control points and adjusting their r value. In 2D, new branches are added by dividing an existing branch into two parts, and joining the two new branches with a third branch at a single control point. The user interface provides feedback when one of the constraints defined above is violated (although constraints on ∇r are enforced by the construction of the model, the non-singularity constraints must be checked after each modification of the model). The method often succeeds even if initialized with a template model that violates the non-singularity constraints. An example of a manually built 3D template model is shown in Fig. 5, and the accompanying movie `manual.mov` demonstrates model editing in 3D.

Template models are deformed to fit an image in an iterative process that minimizes an energy term in presence of constraints. The energy term consists of an image match component and a prior component. The image match component depends on the application. For gray scale image segmentation we use convolution of the image with the Gaussian derivative kernel. The aperture of the kernel is proportional to the local thickness of the object.

When fitting a model to a binary segmentation in 3D, the image match term is the mean square distance to the boundary of the binary object [9]. The volume overlap measure can be used instead, e.g. when fitting coarse models with few control points to highly detailed binary objects.

Regardless of the type of the image match function used, the image match energy component is integrated over the surface of the implied boundary. Integration is performed by sampling the medial spline at a sufficiently fine level of detail. Presently, the spline is sampled uniformly in parameter space. In 3D, the intersections of the trimming edge curve with the sampled grid are added to the set of regular samples collected on the interior of the medial surface.

The prior energy term favors models with low curvature. It imposes a penalty of the form

$$(|\bar{\mathbf{m}}_{i+1} - \bar{\mathbf{m}}_i| - |\bar{\mathbf{m}}_i - \bar{\mathbf{m}}_{i-1}|)^2 \quad (19)$$

on the neighboring control points. The penalty causes the control polygon to be relatively smooth. B-splines possess a minimum curvature property [11] which relates the smoothness of the control polygon to the smoothness of the spline.

The singularity constraints (12) are implemented as heavily weighted penalty functions and are integrated over the sampling grid. The constraints are computed numerically on the order of the sampling grid. Only the first derivatives of (\mathbf{m}, r) are computed during the deformation process. Discrete sampling makes it possible for small local violations of the constraints to occur. The issues of sampling are addressed in the discussion section.

Energy minimization is performed using the $\mu + \lambda$ evolution strategy [20] in two stages. In the similarity transform stage the template is scaled, rotated and translated to best match the image. During the deformation stage,

small groups of adjacent control points are selected in random order and their values are optimized. Each group is optimized for several hundred iterations of the evolutionary algorithm, and several passes over the groups are made, until deformations cease to be significant.

We chose a global non-deterministic optimization method because the image-driven energy function has many local optima. While there exists no mathematical bound on the number of iterations needed for the optimization to converge, the empirical bound in 3D is on the order of 1000 iterations per group of four control points. Convergence time is discussed in section V. The algorithmic complexity is linear in the number of control points because a change to a group of control points only requires the reevaluation of the energy function over a local portion of the object. The complexity is linear in the number of sample points as well.

V. PRELIMINARY RESULTS

We demonstrate the effectiveness of cm-reps by using them for object representation and automatic image segmentation. In 2D, we deform a model of a vertebra to fit a slice of a CT image. In 3D, we deform a model of a hippocampus to fit a manually segmented MRI of the human brain.

In 2D, we segment a lumbar vertebra in an axial slice of the Visible Human abdominal CT subset [41], which is shown in Fig. 7a. Fig. 7b shows a manually constructed initial model of the vertebra with 10 connected branch curves. Notice that this model intentionally contains illegalities, shown as blue segments of the boundary. The result of automatic segmentation is shown in Fig. 7c. The optimization process corrected the illegalities and fitted the image boundaries closely. Fig. 7d shows a detail from fitted model, where in the top left corner the limitation of the non-singularity constraint 15 on the Jacobian can be seen. The curvature constraint 13 was not enforced during this segmentation.

In 3D, we deform a template model of the hippocampus to fit a binary segmentation of the left and right hippocampi in a magnetic resonance image that is a part of a schizophrenia study. We construct a template model of the right hippocampus by manually adjusting control points. We use as a reference the mean right hippocampal image obtained from the statistical shape model computed by Styner et al. using the SPHARM description of all the right hippocampi in the schizophrenia study [31]. We obtain the template model of the left hippocampus by reflection across the midsagittal plane. The template model of the right hippocampus is shown in Fig. 5.

We perform the similarity transform stage of the optimization using volume overlap with the binary segmentation as the likelihood term. We then compute the distance transform of the binary image and blur it using a 1-voxel Gaussian kernel. We perform the deformation stage of the optimization using the mean squared distance likelihood term. The optimization of a model with an order of 4000 sample points typically converges after 10-15 minutes on a

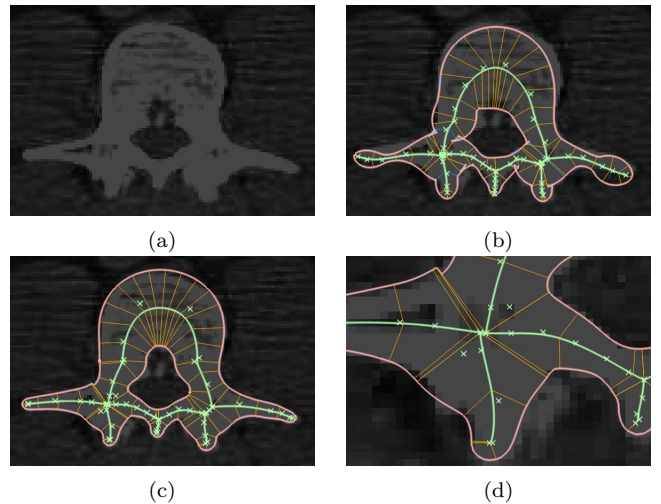


Fig. 7. Automatic segmentation of a vertebra in an abdominal CT image. (a) A slice from the CT image. (b) A manually constructed model with illegalities. The yellow x-shaped marks denote the positions of control points. (c) A result of deforming the model to optimally fit the image. (d) A magnified detail of the model.




Fig. 8. Automatically fitted models of the left and right hippocampi. (a) Left hippocampus: medial surface. (b) Right hippocampus: medial surface. (c) Left hippocampus: implied boundary. (d) Right hippocampus: implied boundary.

900 MHz Intel machine. The resulting models are shown in Fig. 8. Table II shows the mean square distances, in units of voxels, from the boundaries of the hippocampal models to the binary segmentations. Both the right and left hippocampus were fitted accurately. The accompanying movie `deform.mov` demonstrates 3D segmentation.

VI. DISCUSSION AND CONCLUSIONS

The preliminary results show that cm-reps are a promising method for geometric modeling, segmentation and shape representation. In this section we describe the ongoing research (extension to 3D medial branching topology, improved sampling strategy), and compare the method with discrete m-reps.

We use cm-reps to represent 2D objects with a branching

	Left Model	Right Model
Template initialization	216.00	122.59
After similarity transform	4.34	4.55
After deformation	1.71	1.49

TABLE II

MEAN SQUARE DISTANCE FROM LEFT AND RIGHT HIPPOCAMPAL MODELS TO THE BINARY SEGMENTATION.

medial topology and to represent 3D objects with a single medial surface. The extension of the method to cover 3D branching topology remains. The challenge lies in efficiently representing creases formed by seams in medial surfaces and in enforcing constraints along the creases implicitly. B-spline surfaces defined on a quadrilateral mesh can not represent seams efficiently because creases on these surfaces are formed by knot repetition. A crease extends across the entire surface and its image in the parameter space of the b-spline is parallel to the coordinate axes. Recent extensions to the b-spline paradigm [6], [24] allow meshes of arbitrary topology to be used and promise to represent creases more efficiently. In 2D we express the branching constraint (14) in terms of relationships between nearby control points. In 3D, the possibility of expressing the non-linear seam curve constraint (11) in terms of control points remains an open question.

The present implementation uses a uniform grid in b-spline parameter space to sample the cm-rep. Such sampling produces a non-uniform boundary grid, especially near edges of branches, where the medial-to-boundary mapping asymptotes. An adaptive sampling scheme based on distances along the boundary would improve the robustness of segmentation and constraint enforcement.

Since the non-singularity constraints are checked only at the samples, small creases on the boundary are not always detected. We are investigating methods that would allow analytic verification of these constraints on the entire b-spline surface. Alternatively, a search for the maxima of the constraint function, although costly, would detect singularities independently of the sampling.

We now relate continuous m-reps to their discrete cousins. The two representations differ in the strictness of conformation to medial geometry. Continuous m-reps maintain a strict SAT relationship between the boundary and the medial locus. Discrete m-reps allow interpolation of both the boundary and the medial locus but only enforce the SAT relationship at the medial atoms. Differences in stability between the two methods remain to be analyzed.

The two methods differ in the way they handle objects with branching medial topology. Continuous m-reps approximate the behavior of medial loci at branching points. Discrete m-reps do not attempt to simulate the way that medial branches seam and objects with multiple figures are represented by blending the boundaries of single figure m-reps. The blending makes it easy to model and deform complex 3D objects because a figure can be moved relative to another without recomputing the representation. Con-

structive solid geometry is possible with discrete m-reps.

Both representations provide an object-intrinsic metric on the space inside and around the object. This metric allows distance measures to take into account the local thickness of the object. For example, the aperture of the image intensity filter used in grayscale segmentation is proportional to the local thickness of the object. The intrinsic metric can be used to establish an object-centric coordinate system, which provides a framework for fine-scale boundary perturbations within tolerances prescribed by the local thickness [17].

Recent advances in deformable segmentation, registration and shape analysis based on discrete m-reps can readily be extended to cm-reps. Applications that require arbitrary sampling of medial loci can benefit from the continuous representation. For instance, boundary sampling may be adjusted adaptively during segmentation, improving robustness and efficiency. Statistical shape analysis can benefit from cm-reps because correspondences between objects can be represented as a continuous mapping.

In summary, cm-reps are a promising new representation for object modeling, automatic segmentation and shape analysis. Geometric properties of cm-reps are simple and attractive and the legality constraints fit well into the b-spline framework. The preliminary segmentation results are positive and encouraging. The challenges that remain include the extension to cover 3D objects with branching medial topology and application to statistical image analysis.

ACKNOWLEDGEMENTS

The authors thank many UNC MIDAG members for their invaluable advice and support. This work would not have been possible without the inspiration and advice of Prof. Damon. We extend thanks to Prof. Gerig, Dr. Martin Styner and Sean Ho for their advice as well as for the brain MRI they kindly provided us with. This work was carried out under the partial support of NIH grant P01 CA47982. Paul Yushkevich is supported by the Alumni Fellowship of the UNC Department of Computer Science. Some of the equipment was provided under a gift from the Intel Corporation.

REFERENCES

- [1] J. August, K. Siddiqi, and S. Zucker. Ligature instabilities in the perceptual organization of shape. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 1999.
- [2] Robert Blanding, Cole Brooking, Mark Ganter, and Duane Storti. A skeletal-based solid editor. In Willem F. Bronsvort and David C. Anderson, editors, *Proceedings of the Fifth Symposium on Solid Modeling and Applications (SSMA-99)*, pages 141–150, New York, June 9–11 1999. ACM Press.
- [3] Jules Bloomenthal and Ken Shoemake. Convolution surfaces. *Computer Graphics (SIGGRAPH '91 Proceedings)*, 25(4):251–256, July 1991.
- [4] H. Blum and R.N. Nagel. Shape description using weighted symmetric axis features. *Pattern Recognition*, 10(3):167–180, 1978.
- [5] Jonathan C. Carr, Tim J. Mitchell, Richard K. Beatson, Jon B. Cherrie, W. Richard Fright, Bruce C. McCallum, and Tim R. Evans. Reconstruction and representation of 3D objects with radial basis functions. In Stephen Spencer, editor, *Proceedings of the Annual Computer Graphics Conference (SIGGRAPH-01)*, pages 67–76, New York, August 12–17 2001. ACM Press.

- [6] E. Catmull and J. Clark. Recursively generated b-spline surfaces on arbitrary topological meshes. *Computer-Aided Design*, 10:350–355, September 1978.
- [7] David Chen, Stephen M Pizer, and J Turner Whitted. Using multiscale medial models to guide volume visualization. Technical Report TR99-014, University of North Carolina, Chapel Hill, March 2, 1999.
- [8] James Damon. Smoothness and geometry of boundaries associated to skeletal structures. Preprint, 2002.
- [9] P. Danielsson. Euclidean distance mapping. *Computer Graphics and Image Processing*, 14:227–248, 1980.
- [10] Hervé Delingette. Simplex meshes: a general representation for 3D shape reconstruction. In *Conf. on Computer Vision and Pattern Recognition (CVPR '94)*, June 1994.
- [11] G. Farin. *Curves and Surfaces for CAGD: a practical guide*. Academic Press, Inc, San Diego, CA, fourth edition edition, 1995.
- [12] Marie-Paule Gascuel. An implicit formulation for precise contact modeling between flexible solids. In James T. Kajiya, editor, *Computer Graphics (SIGGRAPH '93 Proceedings)*, volume 27, pages 313–320, August 1993.
- [13] Sean M. Gelston and Debasish Dutta. Boundary surface recovery from skeleton curves and surfaces. *Computer Aided Geometric Design*, 12(1):27–51, 1995. ISSN 0167-8396.
- [14] Peter Giblin and Benjamin Kimia. A formal classification of 3d medial axis points and their local geometry. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 566–573, 2000.
- [15] C. M. Hoffmann and P. J. Vermeer. Validity determination for MAT surface representation. In Glen Mullineux, editor, *Proceedings of the 6th IMA Conference on the Mathematics of Surfaces (IMA-94)*, volume VI of *Mathematics of Surfaces*, pages 249–266, Oxford, September 1996. Clarendon Press.
- [16] Takeo Igarashi, Satoshi Matsuoka, and Hidehiko Tanaka. Teddy: A sketching interface for 3d freeform design. *Proceedings of SIGGRAPH 99*, pages 409–416, August 1999. ISBN 0-20148-560-5. Held in Los Angeles, California.
- [17] S Joshi, S Pizer, PT Fletcher, P Yushkevich, A Thall, and JS Marron. Multi-scale deformable model segmentation and statistical shape analysis using medial descriptions. *Invited submission to IEEE-TMI*, page t.b.d., 2002.
- [18] B. Kimia, A. Tannenbaum, and S. Zucker. Shape, shocks, and deformations i: The components of two-dimensional shape and the reaction-diffusion space. *International Journal of Computer Vision*, 15:189–224, 1995.
- [19] Lee Markosian, Jonathan M. Cohen, Thomas Crulli, and John Hughes. Skin: a constructive approach to modeling free-form shapes. *Computer Graphics*, 33(Annual Conference Series):393–400, 1999.
- [20] Zbigniew Michalewicz. *Genetic Algorithms + Data Structures = Evolution Programs*. Springer, Berlin, 1992.
- [21] L. R. Nackman. *Three-Dimensional Shape Description Using the Symetric Axis Transform*. Ph. d. thesis, University of North Carolina at Chapel Hill, 1982. Department of Computer Science.
- [22] M. Näf, O. Kübler, R. Kikinis, M. Shenton, and G. Székely. Characterization and recognition of 3d organ shape in medical image analysis using skeletonization. In *Workshop on Mathematical Methods in Biomedical Image Analysis*, pages 139–150. IEEE Computer Society, 1996.
- [23] R. L. Ogniewicz and O. Kübler. Hierarchic Voronoi skeletons. *Pattern Recognition*, 28(3):343–359, 1995.
- [24] Jörg Peters. Curvature continuous spline surfaces over irregular meshes. *Computer Aided Geometric Design*, 13(2):101–131, 1996.
- [25] L. Piegl and W. Tiller. *The NURBS Book*. Springer-Verlag, New York, New York., 2 edition, 1996.
- [26] S. Pizer, D. Fritsch, P. Yushkevich, V. Johnson, and E. Chaney. Segmentation, registration, and measurement of shape variation via image object shape. *IEEE Transactions on Medical Imaging*, 18:851–865, October 1999.
- [27] Andrei Sherstyuk. Shape design using convolution surfaces. In *Proceedings of Shape Modeling International '99*, March 1999.
- [28] Kaleem Siddiqi, Sylvain Bouix, Allen Tannenbaum, and Steven W. Zucker. The hamilton-jacobi skeleton. In *Proc. Computer Vision*, volume 2, pages 828–834. IEEE, 1999.
- [29] L.H. Staib and J.S. Duncan. Boundary finding with parametrically deformable models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(11):1061–1075, November 1992.
- [30] Duane W. Storti, George M. Turkiyyah, Mark A. Ganter, Chek T. Lim, and Derek M. Stat. Skeleton-based modeling operations on solids. In *SMA '97: Proceedings of the Fourth Symposium on Solid Modeling and Applications*, pages 141–154. ACM, May 1997. held May 14-16, 1997 in Atlanta, Georgia.
- [31] M. Styner and G. Gerig. Medial models incorporating object variability for 3d shape analysis. In *International Conference on Information Processing in Medical Imaging*, pages 503–516, Berlin, Germany, 2001. Springer-Verlag.
- [32] M. Styner. *Combined Boundary-Medial Shape Description of Variable Biological Objects*. PhD thesis, University of North Carolina at Chapel Hill, Chapel Hill, NC, 2001.
- [33] G. Székely, A. Kelemen, Ch. Brechbühler, and G. Gerig. Segmentation of 2-D and 3-D objects from MRI volume data using constrained elastic deformations of flexible Fourier contour and surface models. *Medical Image Analysis*, 1(1):19–34, 1996.
- [34] Pizer S, S Joshi, PT Fletcher, M Styner, G Tracton, and Z Chen. Segmentation of single-figure objects by deformable m-reps. In WJ Niessen and MA Viergever, editors, *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 2208, pages 862–871. Springer, October 2001.
- [35] Ralph Costa Teixeira. *Curvature Motions, Medial Axes and Distance Transforms*. PhD thesis, Harvard University, Cambridge, Massachusetts, June 1998.
- [36] Andrew Thall, P. Thomas Fletcher, and Stephen M. Pizer. Deformable solid modeling using sampled medial surfaces: a multi-scale approach. Technical report TR00-005, University of North Carolina, Chapel Hill, 2000.
- [37] Yea-Fu Tsao and King-Sun Fu. Stochastic skeleton modeling of objects. *Computer Vision, Graphics, and Image Processing*, 1984, 25:348–70, 1984.
- [38] P. J. Vermeer. *Medial Axis Transform to Boundary Representation Conversion*. Ph. d. thesis, Purdue University, 1994.
- [39] P. Yushkevich, Pizer S.M., S. Joshi, and Marron J.S. Intuitive, localized analysis of shape variability. In *International Conference on Information Processing in Medical Imaging*, pages 402–408, Berlin, Germany, 2001. Springer-Verlag.
- [40] S. Zhu and A. Yuille. FORMS: A flexible object recognition and modeling system. *International Journal of Computer Vision*, 20(3):187–212, 1996.
- [41] Visible Human Project. Visible human data set. Technical report, U.S. National Library of Medicine, http://www.nlm.nih.gov/research/visible/visible_human.html, 2001.