Published in final edited form as:

Inf Process Med Imaging. 2019 June ; 11492: 208–220. doi:10.1007/978-3-030-20351-1_16.

Diffeomorphic Medial Modeling

Paul A. Yushkevich¹, Ahmed Aly¹, Jiancong Wang¹, Long Xie¹, Robert C. Gorman², Laurent Younes³, Alison M. Pouch¹

¹Department of Radiology, University of Pennsylvania Perelman School of Medicine, Philadelphia, USA

²Department of Surgery, University of Pennsylvania Perelman School of Medicine, Philadelphia, USA

³Department of Applied Mathematics and Statistics, Johns Hopkins University, Baltimore, USA

Abstract

Deformable shape modeling approaches that describe objects in terms of their *medial axis* geometry (e.g., *m-reps* [10]) yield rich geometrical features that can be useful for analyzing the shape of sheet-like biological structures, such as the myocardium. We present a novel shape analysis approach that combines the benefits of medial shape modeling and diffeomorphometry. Our algorithm is formulated as a problem of matching shapes using diffeomorphic flows under constraints that approximately preserve medial axis geometry during deformation. As the result, correspondence between the medial axes of similar shapes is maintained. The approach is evaluated in the context of modeling the shape of the left ventricular wall from 3D echocardiography images.

1 Introduction

In medical image analysis, *deformable medial models* (or *medial representations*) [5, 10, 13] are frequently used to characterize the shape of anatomical structures that have sheet-like geometry, such as the cerebral cortex, the myocardium, or knee cartilage. They are deformable models that directly characterize the *medial axes* of objects. The medial axis (or *skeleton*) of an object is a set of manifolds formed by all points inside of the object that are equidistant to two or more points on its boundary [3]. The distance from the medial axis to the boundary describes the local thickness of objects. For sheet-like anatomical objects, the shape of the medial axis can be a useful proxy for overall shape, and local thickness is frequently an important biological measurement.

Medial axes of 3D objects can be easily derived using various skeletonization approaches (e.g., [9]). However, even small perturbations of the boundary of an object can result in a significant reconfiguration of the medial axis. This makes it difficult to use medial axes derived by skeletonization in statistical shape analysis. Deformable medial modeling techniques such as *m-reps* [10] overcome this challenge by explicitly controlling the

configuration of the medial axis of the model during deformation (i.e., keeping the number and connectivity of the surfaces in the medial axis fixed). For many classes of anatomical objects such models can approximate the shape of individual objects with good accuracy, while allowing the features derived from the medial axis to be compared across subjects. Deformable medial models have been used to perform statistical shape analysis for various sheet-like brain structures, myocardium and heart valves, abdominal organs, etc. Medial models have also been used to impose geometrical constraints on automatic segmentation of sheet-like structures, e.g., imposing prior knowledge about heart wall thickness during myocardium segmentation [11]. Medial modeling methods include *m-reps* [10] and *s-reps* [5] (which approximate medial axis surfaces using discrete primitives), *cm-reps* [12] (which use splines and subdivision surfaces to model medial axis surfaces), and *boundary-constrained m-reps* [13] (which implicitly model medial geometry by imposing symmetry constraints on a boundary-based model).

This paper addresses a significant limitation of existing medial modeling approaches: they lack a built-in mechanism to prevent models from folding or self-intersecting during deformation; and they do not provide a natural way to extrapolate model transformations to deformations of the ambient space. In recent years, the field of statistical shape analysis has embraced an approach based on *flows of diffeomorphisms*. In this “*diffeomorphometry*” approach [8], shapes are expressed in terms of smooth invertible transformations that deform a canonical *template shape* into each shape of interest. Diffeomorphometry provides a concrete way to impose a metric on the space of shapes, based on the amount of deformation required to map one shape into another, and gives rise to rigorous algorithms for shape interpolation, computation of mean shapes, and other statistics. Diffeomorphometry also allows transformations between shapes to be extended naturally to diffeomorphic transformations of the ambient space, e.g., allowing shape correspondences to be extrapolated to image correspondences.

We propose an algorithm that combines the benefits of deformable medial models (access to intuitive geometric features like thickness) and diffeomorphometry (physically motivated transformation model, rigorous mathematical foundation, invertibility). Our algorithm is formulated as a problem of matching shapes using diffeomorphic flows under constraints that approximately preserve medial axis geometry during deformation. Our approach leverages a general framework for incorporating geometric constraints into diffeomorphic shape matching developed by Arguillère et al. [1]. To our knowledge, this is the first paper to integrate diffeomorphometry and deformable medial modeling. Our experiments focus on demonstrating the feasibility of this method in the context of matching synthetic and real-world anatomical shapes.

2 Materials and Methods

In this section, we first briefly introduce *flows of diffeomorphisms* for shape analysis. We then review the basic concepts of *medial geometry* and define a special class of *diffeomorphic flows that preserve medial geometry*. We then frame the problem of diffeomorphic medial modeling for point-set representations of 3D shapes and implement it

using the general *constrained diffeomorphometry* framework developed by Arguillère et al. [1].

2.1 Flows of Diffeomorphisms for Shape Analysis

In diffeomorphometry, collections of shapes (e.g., hippocampi of different individuals) are represented as diffeomorphic transformations of some template shape $\mathcal{O}$. Following [2], diffeomorphic transformations $\{\phi: \mathbb{R}^n \rightarrow \mathbb{R}^n\}$ form a group under composition. *Flows of diffeomorphisms* are generated by the ordinary differential equation:

$$\frac{\partial \phi(x, t)}{\partial t} = v(\phi(x, t), t), \quad \phi(x, 0) = x, \quad t \in [0, 1], \quad (1)$$

where $v: \mathbb{R}^n \times [0, 1] \rightarrow \mathbb{R}^n$ is a time-varying vector field, called the *velocity field*, that satisfies certain smoothness constraints [8]. Let us write ϕ_t as shorthand for $\phi(\cdot, t)$ and v_t for $v(\cdot, t)$. Applying the diffeomorphism ϕ_t to $\mathcal{O}$ yields a new shape $\phi_t \mathcal{O} = \{x \in \mathbb{R}^n : \phi_t^{-1}(x) \in \mathcal{O}\}$. When comparing shapes, computing shape statistics, or performing shape-based segmentation, a common need is to find a deformation of $\mathcal{O}$ that matches some target shape $\mathcal{O}_{\text{trg}}$ (or some target image I_{trg}) at time $t = 1$, while applying the least total deformation to the space $\mathbb{R}^n$. This is formulated by treating vector fields v_t as elements of a reproducing kernel Hilbert space $\mathcal{H}$ with kernel K , and defining the *kinetic energy* of a diffeomorphic flow generated by v as $E_{\text{kin}}[v] = \frac{1}{2} \int_0^1 \|v_t\|_{\mathcal{H}}^2 dt$, where $\|\cdot\|_{\mathcal{H}}$ is the norm in $\mathcal{H}$. Since pairs of shapes might not have exact correspondence, the shape matching problem is often relaxed to that of approximate shape matching by adding a data attachment term $g(\phi_1 \mathcal{O})$ that measures goodness of fit between the deformed template $\phi_1 \mathcal{O}$ and $\mathcal{O}_{\text{trg}}$ (or I_{trg}) to the energy functional:

$$E[v] = \alpha g(\phi_1 \mathcal{O}) + E_{\text{kin}}[v]. \quad (2)$$

Diffeomorphic flows that minimize $E[v]$ are geodesics in the group of diffeomorphisms. This gives rise to a powerful machinery for statistical shape analysis, in which distances between shapes $\mathcal{O}$ and $\mathcal{O}_{\text{trg}}$ are defined as the kinetic energy of the diffeomorphic flow that minimizes (2) [8]. Using this machinery, well-posed algorithms for statistical shape analysis have been developed [8].

2.2 Medial Axis Geometry

The medial axis is the locus formed by the centers of all maximal inscribed balls in an object [3]. Formally, consider a closed object $\mathcal{O} \subset \mathbb{R}^3$ with no holes and a smooth boundary $\mathcal{B}_{\mathcal{O}}$. Let $B_{x,r} = \{y \in \mathbb{R}^3 : |y - x| \leq r\}$ denote a ball with center x and radius r . $B_{x,r}$ is called a *maximal inscribed ball (MIB)* in $\mathcal{O}$ if $B_{x,r} \subset \mathcal{O}$ and there exists no larger ball $B_{x',r'} \subset \mathcal{O}$ that contains $B_{x,r}$. The *medial axis* of $\mathcal{O}$ is the set of all tuples $(x, r) \in \mathbb{R}^3 \times \mathbb{R}^+$ such that $B_{x,r}$ is an MIB in $\mathcal{O}$. We will call the x -component of the medial axis (i.e., set of all MIB centers) the *medial scaffold* of $\mathcal{O}$, and denote it $\mathcal{M}_{\mathcal{O}}$, and we will call the r -component the *radial scalar*

field, denoted $\mathcal{R}_\mathcal{O}$. Generically¹, $\mathcal{M}_\mathcal{O}$ consists of a set of surfaces, called *medial surfaces*. These medial surfaces are bounded by curves that include curve segments that are shared by multiple medial surfaces, called *seam curves*, and curve segments that belong to just one medial surface, called *free edges*. As shown in [4], MIBs centered on the interior of medial surfaces are tangent to $\mathcal{B}_\mathcal{O}$ at two points. MIBs centered on seam curves are tangent to $\mathcal{B}_\mathcal{O}$ at three points, and MIBs centered along free edges are tangent to $\mathcal{B}_\mathcal{O}$ at just one point. Each point in $\mathcal{B}_\mathcal{O}$ belongs to exactly one MIB.

2.3 Medial Structure Preservation Under Diffeomorphic Flow

A flow ϕ acting on an object $\mathcal{O}$ is called *medial-preserving* if $\phi_t \mathcal{M}_\mathcal{O} = \mathcal{M}_{\phi_t \mathcal{O}}$ for all $t \in [0, 1]$, i.e., the medial scaffold of the transformed object coincides with the transformed medial scaffold of the original object. Medial preserving flows maintain point-wise correspondences on the medial scaffold as the object deforms, allowing statistical analysis of properties derived from the medial axis, such as thickness. By contrast, arbitrary diffeomorphic flows that are not medially preserving may cause the medial scaffold to change structure, e.g., by adding new surfaces or reconfiguring their geometry [4]. **The aim of this paper is to perform approximate shape matching using medial-preserving flows.**

Let the tuple of points $\{b_1, \dots, b_n; x\}$ be called a *medial tuple* in $\mathcal{O}$ if there exists an MIB centered on x that is tangent to $\mathcal{B}_\mathcal{O}$ at points $b_1, \dots, b_n$. Let ϕ be a flow of diffeomorphisms that satisfies the following property: for every medial tuple $Z = \{b_1, \dots, b_n; x\}$ in $\mathcal{O}$ and every time $t \in [0, 1]$, the transformed tuple at time t , denoted $\phi_t Z = \{\phi_t(b_1), \dots, \phi_t(b_n); \phi_t(x)\}$, is also a medial tuple in $\phi_t \mathcal{O}$. Then it can be shown (proof omitted due to limited space) that ϕ is medially preserving. Let us denote $\phi_t Z$ as $Z^t = \{b_1^t, \dots, b_n^t; x^t\}$. To check if the tuple Z^t (with $n > 1$) is a medial tuple in $\phi_t \mathcal{O}$ it suffices to check three conditions:

$$(x^t - b_i^t) \perp T_{b_i^t} \quad \forall i \in [1, n] \quad (3)$$

$$|x^t - b_i^t| = |x^t - b_1^t| \quad \forall i \in [2, n] \quad (4)$$

$$|x^t - b'| \geq |x^t - b_1^t| \quad \forall b' \in \phi_t \mathcal{B}_\mathcal{O}, \quad (5)$$

where $T_{b_i^t}$ denotes the tangent plane to the boundary surface $\phi_t \mathcal{B}_\mathcal{O}$ at the point b_i^t . Conditions (3) and (4) ensure that x^t is the center of a ball that is tangent to $\phi_t \mathcal{B}_\mathcal{O}$ at points $b_1, \dots, b_n$, and condition (5) ensures that this ball is inscribed in $\phi_t \mathcal{O}$, and thus an MIB in $\phi_t \mathcal{O}$. Note that conditions (3) and (4) only involve local geometry of $\phi_t \mathcal{B}_\mathcal{O}$ around the tuple, while condition (5) requires the knowledge of the entire $\phi_t \mathcal{B}_\mathcal{O}$. In practice, it appears to be sufficient to only

¹A property of $\mathcal{O}$ is considered “generic” if it is invariant to small smooth perturbations of $\mathcal{O}$. For example, the centers of the MIBs in a perfect cylinder form a line, but a small perturbation breaks this perfect symmetry.

enforce the first two conditions, and we conjecture that diffeomorphic flows that satisfy only (3) and (4) for all Z^t with $n > 1$ are medially preserving.

2.4 Optimal Control Framework for Diffeomorphometry

Above, we defined a set of local geometric constraints that we want flows of diffeomorphisms to satisfy. We now briefly summarize a framework for constrained diffeomorphic shape matching developed by Arguillère et al. in [1]. We restrict our attention to shapes represented by point sets, which is relevant for the implementation of medial constraints in this paper.

Let $q^0 = \{q_1^0, \dots, q_k^0\}$ be a finite set of points in $\mathbb{R}^n$, e.g., a set of landmarks sampled on the boundary (and perhaps on the interior) of the template shape $\mathcal{O}$. Let q^t denote the positions of these points at time t , i.e., $q^t = \phi_t q^0 = \{\phi_t(q_1^0), \dots, \phi_t(q_k^0)\}$. Let us express the data attachment term $g(\phi_1 \mathcal{O})$ purely in terms of these landmarks, i.e., $g(\phi_1 \mathcal{O}) = g(q^1)$. The matching of $\mathcal{O}$ to the target shape/image then takes the form of minimizing the energy

$$E[v] = \alpha g(q^1) + \frac{1}{2} \int_0^1 \|v_t\|_{\mathcal{H}}^2 dt, \quad (6)$$

where ϕ_t is related to v via (1). This is a kind of optimal control problem, in which v is a “control”, i.e., a function that affects the paths of points q^0 and guides them towards their target locations. It is a remarkable fact that this problem can be restated as a problem involving a much simpler set of controls associated with each landmark. As proven in [1], if v is a minimizer of $E[v]$ in (6), then v can be interpolated from a set of vector-valued functions $u(t) = \{u_1(t) : [0, 1] \rightarrow \mathbb{R}^n, \dots, u_k(t) : [0, 1] \rightarrow \mathbb{R}^n\}$ associated with the landmarks. The interpolation is via the kernel K of the Hilbert space $\mathcal{H}$. For every t ,

$$v_t(x) = \sum_{j=1}^k K(x, q_j^t) u^j(t). \quad (7)$$

The function $u_i(t)$ is called the *momentum* of the landmark q_i^0 and can be visualized as a vector placed at each point along the path traced by q_i^0 . In general, the kernel $K(x, y)$ is a matrix-valued function, but in most practical applications, it is chosen to be of the form $K(x, y) = \eta(\|x - y\|) \mathcal{I}_{n \times n}$, where $\eta : \mathbb{R} \rightarrow \mathbb{R}$ is a radial basis kernel. As in [1], we use a Gaussian kernel $\eta(z) = e^{-0.5 z^2 / \sigma^2}$.

Given the interpolation formula (7), $\|v_t\|_{\mathcal{H}}$ and subsequently $E[v]$ can be expressed entirely in terms of the momenta $u(t)$:

$$E[u] = \alpha g(q^1) + \frac{1}{2} \int_0^1 \underbrace{\sum_{i=1}^k \sum_{j=1}^k u_i(t)^T K(q_i^t, q_j^t) u_j(t)}_{\|v_t\|_{\mathcal{H}}^2} dt, \quad (8)$$

where the landmark positions are updated according to:

$$\dot{q}_i^t = v_t(q_i^t) = \sum_{j=1}^k K(q_i^t, q_j^t) u_j(t). \quad (9)$$

Diffeomorphic methods often take advantage of the fact that minimizers of $E[u]$ are geodesics in the group of diffeomorphisms and thus are determined by the initial momentum $u(0)$. However, here we are concerned with minimizing $E[u]$ subject to a set of constraints $C_1, \dots, C_L$, each in the form $C_l(q^t) = 0$. Arguillère et al. [1] propose two strategies for solving such problems: as an initial momentum problem where geodesics are computed by projection of $\dot{q}_i^t$ into the null space of the constraints; and using an *Augmented Lagrangian* (AL) method in which optimization is over the complete momentum $u(t)$ and the constraints are added as penalty terms to $E[u]$. We adopt the AL-based method, which is easier to implement and is numerically more stable, but converges slowly [1].

In the AL method, minimization of $E[u]$ subject to constraints $\{C_l(q^t) = 0\}$ is achieved as a series of unconstrained problems, indexed by $m = 1, \dots, M$. A set of L Lagrange multiplier functions $\lambda_l^m : [0, 1] \rightarrow \mathbb{R}$ corresponding the constraints is introduced, initialized $\lambda_l^1(t) = 0$. At iteration m , the following energy is minimized:

$$E_{AL}^m[u] = E[u] + \int_0^1 \sum_{l=1}^L \left[\frac{\mu^m}{2} C_l(q^t)^2 - \lambda_l^m(t) C_l(q^t) \right] dt, \quad (10)$$

where μ^m is a scalar weight. The Lagrange multipliers and weights are updated at each iteration as

$$\lambda_l^{m+1}(t) = \lambda_l^m(t) - \mu^m C_l(q^t); \quad \mu^{m+1} = \mu^m \cdot \mu,$$

where $\mu = 10$ is a constant scaling factor. Minimization of (10) only requires computing the gradient of $E_{AL}^m[u]$ with respect to $u(t)$, which involves solving an ODE backwards in time, as specified in [1, Remark 24].

2.5 Medial Constraints for Point Set Diffeomorphicmetry

In this paper, we adapt the Arguillère et al. [1] approach to the problem of matching shapes using medial-preserving flows of diffeomorphisms. To reduce the problem to point sets, we relax the constraints (3) and (4) to a finite set of medial tuples $Z_1, \dots, Z_p$ in $\mathcal{O}$, rather than

for the infinite set of all medial tuples in $\mathcal{O}$. We also discretize the time dimension. Although by enforcing constraints (3) and (4) for a discrete set of tuples and timepoints no longer guarantees the exact preservation of medial structure by the diffeomorphic flow, in practice, solving the discrete problem results in near-preservation of medial structure, i.e., $\mathcal{M}_{\phi_t \mathcal{O}} \simeq \phi_t \mathcal{M}_{\mathcal{O}}$, which is sufficient in order to perform statistical shape analysis using point correspondences between $\mathcal{M}_{\mathcal{O}}$ and $\phi_t \mathcal{M}_{\mathcal{O}}$.

Given a template object $\mathcal{O}$, the medial tuples $Z_1, \dots, Z_p$ in $\mathcal{O}$ are sampled “regularly” from among all medial tuples in $\mathcal{O}$. Such sampling is nontrivial, because regular sampling of the x ’s on the medial scaffold results in irregular sampling of b ’s on the boundary, and vice-versa. Since constructing a template is a one-off task, we manually sample medial tuples in the template object $\mathcal{O}$ using a GUI tool. The tool computes the approximate medial scaffold $\mathcal{M}_{\mathcal{O}}$ using pruned Voronoi skeletonization [9]. Samples are taken along the seams and free edges in $\mathcal{M}_{\mathcal{O}}$ as well as on the surfaces forming $\mathcal{M}_{\mathcal{O}}$. Samples are organized into a triangle mesh, which is then inflated of both sides of each triangle to create a boundary triangle mesh. Template-building is shown in Fig. 1(a).

The Arguillère et al. approach [1] is applied as follows. We let q^0 consist of all of the points contained in the tuples $Z_1, \dots, Z_p$. Each tuple Z_i indexes n_i boundary landmarks and one medial landmark. Let $\mathfrak{B}[i, j]$ denote the index of the j -th boundary landmarks in Z_i and let $\mathfrak{M}[i]$ denote the index of the medial point in Z_i . Then, at time t , the transformed i -th medial tuple is given by $Z_i^t = \{q_{\mathfrak{B}[i, 1]}^t, \dots, q_{\mathfrak{B}[i, n_i]}^t, q_{\mathfrak{M}[i]}^t\}$. We formulate a set of constraints that are discrete equivalents of the constraints (3) and (4).

Constraint (3) requires the vector $q_{\mathfrak{B}[i, j]}^t - q_{\mathfrak{M}[i]}^t$ to be orthogonal to $\phi_t \mathcal{B}_{\mathcal{O}}$ at the point $q_{\mathfrak{B}[i, j]}^t$. A pair of vectors spanning the tangent plane to $\phi_t \mathcal{B}_{\mathcal{O}}$ at $q_{\mathfrak{B}[i, j]}^t$ can be approximated as the weighted sum of the coordinates of adjacent boundary vertices, e.g., using the Loop tangent scheme [14, p.71]. We can write these tangent vectors as $W_{i, j}^1 q^t$ and $W_{i, j}^2 q^t$ where $W_{i, j}^1$ and $W_{i, j}^2$ are sparse $3 \times k$ matrices. The discrete equivalents of constraints (3) and (4) are then:

$$\begin{aligned} (q_{\mathfrak{B}[i, j]}^t - q_{\mathfrak{M}[i]}^t)^T W_{i, j}^\gamma q^t &= 0 \quad \forall j \in [1, n_i], \gamma = 1, 2, \\ |q_{\mathfrak{B}[i, j]}^t - q_{\mathfrak{M}[i]}^t|^2 - |q_{\mathfrak{B}[i, j]}^t - q_{\mathfrak{M}[i]}^t|^2 &= 0 \quad \forall j \in [2, n_i]. \end{aligned} \quad (11)$$

Following Pizer et al. [10], we refer to the vectors $q_{\mathfrak{B}[i, j]}^t - q_{\mathfrak{M}[i]}^t$ as *spokes* (as in spokes on a bicycle wheel). The constraints state that spokes must be orthogonal to the boundary and that the lengths of the spokes must be equal. When monitoring AL optimization, we describe the extent to which the constraints are violated in terms of two related intuitive metrics: *spoke-normal deviation*, which measures the angle between the spoke $q_{\mathfrak{B}[i, j]}^t - q_{\mathfrak{M}[i]}^t$ and the boundary normal vector at $q_{\mathfrak{B}[i, j]}^t$; and *spoke-spoke mismatch*, which measures the maximum relative error between the lengths of two spokes in the tuple Z_i .

2.6 Data Attachment Term

Various data attachment terms can be used. When the target locations q_i^{trg} of the boundary and/or medial landmarks are given, a simple sum of square distances (SSD) attachment term is used: $g(q^1) = \sum_{i \in Y} |q_i^1 - q_i^{\text{trg}}|^2$, where Y denotes the subset of landmarks with given target locations. When fitting a model to a target binary image $\mathcal{I}_{\text{trg}}$, we define $g(q^1)$ as an approximation of the Dice similarity coefficient (DSC) between the interior of $\phi_1\mathcal{O}$ and the foreground region of the target binary image. The DSC attachment term leverages the triangulation of the boundary and medial surfaces created when medial tuples are sampled. Points in a 3D wedge formed by each medial triangle and the corresponding boundary triangle are sampled regularly, and the image $\mathcal{I}_{\text{trg}}$ is sampled at each sample point. Each sample point is also assigned a volume element. Integrating the sampled intensity values times the volume element yields an approximation of overlap volume, while the volume of $\phi_1\mathcal{O}$ is computed by just integrating the volume element. The volume of the foreground region in $\mathcal{I}_{\text{trg}}$ is known a priori. This allows DSC and its gradient with respect to q^1 to be computed.

2.7 Numerical Implementation

The AL minimization is implemented in C++ using CPU multi-threading. ODEs to compute $E_{\text{AL}}^m[u]$ and its gradient are solved numerically using Euler's method with 40 time steps.

Unconstrained minimization (10) at each AL iteration is performed using a pseudo-Newton method (LBFGS) and is allowed to proceed until the change in $E_{\text{AL}}[u]$ (or $E_{\text{AL}}[u, \eta, \rho]$) falls below a tolerance threshold, usually for 1000–4000 iterations. Constraints, which are a mix of quadratic and linear expressions, are implemented as sparse matrix-vector multiplication.

The scaling factor μ in the AL algorithm is set to 10, while the initial value of μ is set experimentally, as is the data attachment term weight, α and the standard deviation σ of the Gaussian kernel.

3 Experiments and Results

We perform experiments on synthetic and real-world anatomical shapes, with the focus on demonstrating the feasibility of matching templates to target shapes using medial-preserving diffeomorphic flows.

3.1 Synthetic Shape Example

We use a toy example to show that our method can match shapes with branching medial axes, a problem that not all medial modeling approaches can solve (e.g., not [12]). We created two 3D shapes that have a medial axis consisting of three surfaces joining along a single seam curve, as follows. First, we manually painted a binary 3D object and smoothed it to create a surface labeled “source shape” in Fig. 1. We then flipped the binary image about the x -axis and performed diffeomorphic deformable image registration between the original image and the flipped image. The resulting warp was applied to the source shape, yielding

the shape labeled as “target shape” in Fig. 1. By using deformable registration, we obtained point-wise correspondences between the two shapes.

To create the initial template medial model, we used a GUI tool as described in Sect. 2.5. The Voronoi skeleton of the “source shape” (after additional smoothing) is shown in Fig. 1(a) along with the medial points sampled from the skeleton and the triangulation of the medial points. The medial mesh shown in Fig. 1(a) was inflated and subdivided once using the Loop scheme [7] and then fitted to the “source shape” in Fig. 1 using the deformable medial modeling method in [13]. This yielded an initial template for which the medial constraints (11) are fully satisfied. The boundary and the medial model of this template are visualized in Fig. 1(b) under $t = 0$.

We then fitted the template to the target shape using our medially-constrained diffeomorphometry approach. We initialized the fitting using the SSD data attachment term, with target landmark locations obtained by applying the image registration warp to the template’s boundary landmarks. This was done using a single AL iteration with $\mu = 1$, since the goal was only to initialize the subsequent DSC-based optimization. We then performed full AL optimization with the DSC data attachment term, and with μ increasing from 0.001 to 10 over several AL iterations. Figure 2 plots the change in the data attachment, kinetic energy, total constraint violation, spoke-normal deviation, and spoke-spoke mismatch over the course of optimization. As μ increases (>0.1), the data attachment term begins to retreat from its highest values in order to satisfy the constraints. However, the resulting practical gain at higher values of μ is very small, as the maximal normal-spoke deviation and spoke-spoke mismatch are already quite small for $\mu = 0.1$: about 1° and 1%, respectively, and don’t decrease dramatically as μ increases. The flow visualized in Fig. 1(b) corresponds to the solution with $\mu = 0.1$, which may be considered the “sweet spot” of the optimization.

The three rows of shapes in Fig. 1(b) correspond to the evolving template boundary ($\phi_t \mathcal{B}_0$), the evolving template medial scaffold ($\phi_t \mathcal{M}_0$), and the Voronoi skeleton of the evolving template (an approximation of $\mathcal{M}_{\phi_t \mathcal{O}}$). Critically, $\phi_t \mathcal{M}_0$ and $\mathcal{M}_{\phi_t \mathcal{O}}$ are highly consistent, indicating that our *medial preservation constraints are working*. By contrast, Fig. 1(c) shows that *conventional point set diffeomorphometry does not preserve medial axis structure*. We performed point set geodesic shooting between the boundary of our template ($\mathcal{B}_0$) and the boundary of the fitted template ($\mathcal{B}_{\phi_1 \mathcal{O}}$) and extracted the Voronoi skeleton of the evolving shape. As pointed out by the yellow arrows in Fig. 1(c), the skeleton undergoes structural change during deformation.

3.2 3D Echocardiography Data

To demonstrate the feasibility of diffeomorphic medial modeling on real-world imaging datasets, we use our approach to model the shape of the left ventricle (LV) wall in a dataset of 14 transesophageal 3D echocardiograms acquired prior to heart valve surgery. The dataset includes images from 8 patients who underwent surgery for ischemic mitral regurgitation (IMR) and 6 patients who did not have surgery. The LV wall was manually segmented at systole in each case by trained raters. The first IMR case was used to construct a template medial model using the same approach as in the toy example. Affine registration guided by

five manually-placed landmarks was performed between the binary segmentation images of the first IMR case and all other cases. This was needed to correctly line up the LV outflow tract between cases. To initialize our medially-constrained diffeomorphometry approach with the affine registration parameters, we first performed one iteration of AL optimization with the SSD data attachment term, where the affine-transformed template landmarks were treated as target locations q^{reg} . We then performed multiple iterations of AL optimization with the DSC data attachment term. Figure 3 plots the DSC data attachment term, kinetic energy, total constraint penalty, maximum spoke-normal deviation and maximum spoke-spoke mismatch over the course of AL optimization. Fitting performance was highly consistent across the 14 cases, with final Dice coefficient ranging between 0.935 and 0.945 for all cases, and the maximum normal-spoke deviation and spoke-spoke mismatch having ranges 0.04° – 1.08° and 0.1% – 2.2% respectively. Overall, this demonstrates excellent ability of medial-preserving diffeomorphic flows to fit real-world anatomical shapes. Figure 4 shows an example of the LV template deforming to match an individual case. As in the toy example, the preservation of the template's medial scaffold is demonstrated by excellent consistency between the deforming medial scaffold of the template $(\phi_t \mathcal{M}_\phi)$ and the Voronoi skeleton of the deforming template $(\mathcal{M}_{\phi_t \phi})$.

4 Discussion and Conclusions

To our knowledge the method presented here is the first attempt to directly combine medial modeling and diffeomorphometry (an indirect approach to combining cm-reps and diffeomorphometry was recently proposed by Hong et al. [6]). The current paper demonstrates the feasibility of shape matching with this approach, but additional work is needed to extend all the tools of computational anatomy [8] (e.g., computing mean shapes or principal geodesics) to diffeomorphic medial modeling. A limitation of our approach is its high computational burden, with thousands of iterations needed for AL optimization to converge and each iteration having cost quadratic in the number of landmarks. The heart shape experiments were performed in an hour per case on a 16-core CPU. An open-source C++ implementation of our method is at <https://github.com/pyushkevich/cmrep>.

Acknowledgments

This work is supported by NIH grants AG056014, EB017255, HL103723, HL141643.

References

1. Arguillère S, Trélat E, Trounev A, Younes L: Shape deformation analysis from the optimal control viewpoint. *Journal de mathématiques pures et appliquées* 104(1), 139–178 (2015)
2. Arnold VI: Sur la géométrie différentielle des groupes de lie de dimension infinie et ses applications à l'hydrodynamique des fluides parfaits. *Ann. Inst. Fourier* 16(1), 319–361 (1966)
3. Blum H, Nagel RN: Shape description using weighted symmetric axis features. *Pattern Recogn.* 10(3), 167–180 (1978)
4. Giblin PJ, Kimia BB: On the local form and transitions of symmetry sets, medial axes, and shocks. *Int. J. Comput. Vis.* 54(1–3), 143–157 (2003)
5. Hong J, Vicory J, Schulz J, Styner M, Marron JS, Pizer SM: Non-euclidean classification of medically imaged objects via S-reps. *Med. Image Anal.* 31, 37–45 (2016) [PubMed: 26963609]

6. Hong S, Fishbaugh J, Gerig G: 4D continuous medial representation by geodesic shape regression. In: 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018), pp. 1014–1017. IEEE (2018)
7. Loop C, DeRose T: Generalized B-spline surfaces of arbitrary topology. In: Computer Graphics (ACM SIGGRAPH Proceedings), pp. 347–356 (1990)
8. Miller MI, Trounev A, Younes L: Geodesic shooting for computational anatomy. *J. Math. Imaging Vis* 24(2), 209–228 (2006) [PubMed: 20613972]
9. Näf M, Kübler O, Kikinis R, Shenton M, Székely G: Characterization and recognition of 3D organ shape in medical image analysis using skeletonization. In: Workshop on Mathematical Methods in Biomedical Image Analysis, pp. 139–150. IEEE Computer Society (1996)
10. Pizer SM, et al.: Deformable m-reps for 3D medical image segmentation. *Int. J. Comput. Vis.* 55(2), 85–106 (2003) [PubMed: 23825898]
11. Sun H, Frangi AF, Wang H, Sukno FM, Tobon-Gomez C, Yushkevich PA: Automatic cardiac MRI segmentation using a biventricular deformable medial model. *Med. Image Comput. Comput. Assist. Interv* 13(Pt 1), 468–475 (2010) [PubMed: 20879264]
12. Yushkevich PA, Zhang H, Gee J: Continuous medial representation for anatomical structures. *IEEE Trans. Med. Imaging* 25(2), 1547–1564 (2006) [PubMed: 17167991]
13. Yushkevich PA, Zhang HG: Deformable modeling using a 3D boundary representation with quadratic constraints on the branching structure of the blum skeleton. *Inf. Process. Med. Imaging* 23, 280–291 (2013) [PubMed: 24683976]
14. Zorin D, Schröder P, DeRose T, Kobbelt L, Levin A, Sweldens W: Subdivision for modeling and animation In: SIGGRAPH Course Notes. ACM, New York (2000)

Fig. 1.

Template generation and shape matching in the toy example. (a) Manual template construction by sampling tuples from the skeleton of the source shape. (b) Medial-preserving diffeomorphic flow between the template and a target shape. The deforming medial scaffold of the template ($\phi_t \mathcal{M}_O$) is in close agreement with the Voronoi skeleton of the deforming template ($\mathcal{M}_{\phi_t O}$), which illustrates that medial preservation constraints are working. (c) When analogous same shape matching (between $\mathcal{B}_O$ and $\mathcal{B}_{\phi_1 O}$) is performed using point set geodesic shooting *without medial preservation constraints*, the Voronoi skeleton of the deforming object changes structure and bifurcates (yellow arrows)

Fig. 2. Data attachment, kinetic energy, and constraints over the course of AL optimization for synthetic shape matching.

Fig. 3. Data attachment, kinetic energy, and constraints over the course of AL optimization for the 14 LV wall matching cases.

Fig. 4.

Example of shape matching in the LV wall dataset. The top row shows the deforming medial scaffold (colored by thickness = $2 \times$ MIB radius) and the deforming boundary as a white wireframe. The bottom row shows the Voronoi skeleton of the deforming boundary, again emphasizing the preservation of medial scaffold structure during deformation.