

Population-based Fitting of Medial Shape Models with Correspondence Optimization

Timothy B. Terriberry¹, James N. Damon², Stephen M. Pizer¹, Sarang C. Joshi⁴, and Guido Gerig^{1,3}

¹ Dept. of Computer Science, Univ. of North Carolina, Chapel Hill, NC 27599, USA

² Dept. of Mathematics, Univ. of North Carolina, Chapel Hill, NC 27599, USA

³ Dept. of Psychiatry, Univ. of North Carolina, Chapel Hill, NC 27599, USA

⁴ Dept. of Biomedical Engineering, Univ. of Utah, Salt Lake City, UT 84112, USA

{tterribe,pizer,gerig}@cs.unc.edu, jndamon@math.unc.edu,
sjoshi@sci.utah.edu

Abstract. A crucial problem in statistical shape analysis is establishing the *correspondence* of shape features across a population. While many solutions are easy to express using boundary representations, this has been a considerable challenge for medial representations. This paper uses a new 3-D medial model that allows continuous interpolation of the medial manifold and provides a map back and forth between it and the boundary. A measure defined on the medial surface then allows one to write integrals over the boundary and the object interior in medial coordinates, enabling the expression of important object properties in an object-relative coordinate system. We use these integrals to optimize correspondence during model construction, reducing variability due to the model parameterization that could potentially mask true shape change effects. Discrimination and hypothesis testing of populations of shapes are expected to benefit, potentially resulting in improved significance of shape differences between populations even with a smaller sample size.

1 Introduction

In questions of statistical shape analysis, the foremost is how such shapes should be represented. The number of parameters required for a given accuracy and the types of deformation they can express directly influence the quality and type of statistical inferences one can make. Most methods of establishing correspondence in a population use features on the boundary [1–4], since in many imaging modalities the interior of objects have a uniform appearance with poorly localized features. However our research uses a medial model parametrization, which represents a solid object using a skeleton of a lower dimension and naturally expresses intuitive changes such as “bending”, “twisting”, and “thickening”, but where establishing correspondence is more difficult. As a descriptor of *shape*, the medial axis can be used to provide a detailed quantitative and qualitative analysis that simpler object descriptors, such as volume, surface area, pose, etc., cannot. Pizer et al. give an overview and comparison of definitions and numerous methods for computing of a medial axis [5]. Yet the reversal of the original




Fig. 1. A dense sampling of a medial model of a left caudate defined continuously using subdivision surfaces, with volume overlap $V_{\text{Dice}} = 93.46\%$.

relationship, from an object describing a medial axis to a medial axis describing an object, is the critical idea that makes medial representations effective image analysis tools. It replaces an inherently unstable and ill-defined problem—computing the medial axis of an object from its boundary—with a well-defined and stable one: computing the boundary of an object from its medial axis. An object is then modeled by deforming a template medial axis until the associated boundary matches that of the target object.

Näf et al. use the medial axis to measure local bone thickness and to describe the sulco-gyral foldings of the human brain [6]. Zhang et al. use the medial axis for articulated shape matching [7]. They use Siddiqi et al.’s shock detection algorithm to construct the medial axis and classify the voxels according to Giblin and Kimia’s taxonomy directly via [8]. Golland et al. use “fixed-topology skeletons” for 2D shape classification [9, 10]. Similar to Golland’s fixed-topology skeletons, Pizer et al. introduce a sampled medial representation called (discrete) *m-reps* used for segmentation [11, 12]. Bouix et al. apply the medial axis to estimate the local width of the hippocampus [13]. They propose two different approaches to solving the problem of identifying corresponding points between different subjects: projecting the axis onto a plane and rigidly aligning it or using nonlinear deformations to warp the axis to a common template shape. Styner et al. [14] use sampled medial descriptions for shape modeling and analysis in clinical studies of hippocampi and lateral ventricles. They use boundary correspondence established via spherical harmonics (SPHARMs) when initializing the medial model fitting process, but do not enforce it during the individual optimization for each subject, allowing features to wander independently.

Statistical analysis of populations requires appropriate solutions for *robust parameterizations* of shape models and *establishing correspondence* across a class of objects that is meaningful given the specific task and application domain. An inherent problem in any shape representation, one must ensure that the parameters of the representation in some sense control the “same” features of

the resulting shape. Otherwise, noise in the parameterization can overwhelm the size of any shape change effect, reducing or eliminating the power of the tests.

The method presented here does *not* provide yet another means for determining the correspondence between shapes. Instead, it takes the stance that correspondence is inherently specific to the problem domain. Without some input governing the process whereby the representation is constructed results are unpredictable. Our modeling scheme can use any existing correspondence on the boundary, defined via prior or learned knowledge of the target population, to establish correspondence on the medial representation. This in turn requires a unique link between locations on the boundary and on the medial parametrization. In principle, the medial geometry provides an intrinsic link between the boundary and the medial axis that allows one to use the appropriate representation for the desired analysis task. Every (non-singular) point on the medial axis lies at the center of ball which is tangent to the boundary in two places, and the field of vectors from the points on the axis to their associated points of tangency, called *spokes*, provides this link. This relationship was explored in [15], and [16, 17] provide an analysis of the differential geometry in arbitrary dimension. However with the discrete version of m-rep models [11], the connection with the boundary is given only at a coarse set of discrete points and the interpolation given by [18] to recover a dense sampling does not respect the intrinsic medial geometry. Correspondence is established only approximately using a regularizing term in the model fit optimization, which requires determining an appropriate weight and trades off homology for goodness of fit. More recent work on interpolation does respect the intrinsic medial geometry [19], but it sacrifices the uniqueness of the representation, requires expensive numeric integration, and only approximately interpolates the original model.

An alternative approach to discrete skeletons, *cm-reps*, considers the dual problems of designing a discrete computer representation and computing a continuous mathematical representation from it as a coupled system [20–23]. A continuous representation inherently provides the necessary connection between the medial axis and the boundary and thus is chosen for this paper.

2 Group-wise model fitting with explicit correspondence optimization

This paper addresses the problem of producing a collection of models that represent a set of shapes to be used in statistical tests. The basic approach is to start with a template object, described medially with a fixed branching topology, to align a copy of that template to each target shape, and then to deform it to match. This work only considers fitting a model to an existing, segmented shape, such as a binary image or a triangulated surface.

The following sections begin by describing a method for sampling a continuous medial axis and then show how this sampling can be used to approximate medial integrals. The complete process for fitting a single model to a single target shape follows. This involves aligning a template to the target shape and deform-




Fig. 2. The entire population of objects is mapped to a common coordinate system to optimize correspondence.

ing it using a multi-scale constrained optimization. Finally, the process for a single shape is extended to a technique for producing a population of models with a common correspondence. This is integrated directly into the deformation stage by constructing deformations approximately transverse to the fitting process. The models are mapped to a common coordinate system and the control points are adjusted to achieve a parameterization that matches an explicit correspondence given on the boundary, as illustrated in Figure 2. This is done by taking advantage of the inherent link between the boundary and the medial axis given by a continuous medial model, moving the problem of formulating correspondence to the boundary, where most of the information lies.

The specific discretized representation for the target shapes used here are binary images, $I(x) : \mathbb{R}^3 \rightarrow \{0, 1\}$, defined to be 1 when x is in the shape’s interior, and 0 elsewhere. Other representations like triangle meshes could be used as well, as described in [24].

2.1 Approximating Medial Integrals

The main mathematical tool used in this paper is the medial integral, which we approximate via numeric integration. This section describes the procedure.

Sampling the Medial Axis. The first step is to define a sampling of the medial axis. We use the subdivision surface approach presented in [23] for our 3D continuous medial axis representation, illustrated in Fig. 1. This divides the medial axis into a fixed number of patches which can be evaluated analytically at arbitrary points using B-splines or other fast evaluation methods [25]. Each patch gives a continuous function m , the position of a point on the axis, and r , the radius of the maximally inscribed sphere, defined over a square domain $(u, v) \in [0, 1]^2$. The exact expressions for m and r and their derivatives can be found in [24] and are omitted for space reasons. The unit spoke vectors U^\pm pointing towards the two points of tangency on the boundary and the two boundary points themselves are computed directly from derivatives of m and r by

$$U^\pm = -\nabla r \pm \sqrt{1 - \|\nabla r\|^2} \cdot \mathcal{N}, \quad \mathcal{B}^\pm = m + rU^\pm, \quad (1)$$

where ∇r is the Riemannian gradient and \mathcal{N} is the unit normal vector to the medial surface. A single sample is then placed at the center of each patch with extent $(\Delta u, \Delta v) = (1, 1)$ and recursively subdivided until

$$\|\mathcal{B}_u^\pm \Delta u \times \mathcal{B}_v^\pm \Delta v\| < \tau^2 \quad (2)$$

for some threshold τ , which ensures the sample area of the boundary is sufficiently small on both sides of the medial axis. A sample set \mathbb{S} is then constructed from these samples, containing the tuple (m, r, U) and all its derivatives, where U is one of the two spoke vectors U^\pm . Each point on the medial axis contributes two sample values, one for each spoke vector, in order to integrate over both sides of the medial axis.

Numeric Integration in Medial Coordinates. Damon showed how to rewrite volume and surface integrals of a medially defined region in terms of medial integrals [26]. We begin with a motivational example: some simple volume integrals over the object interior Ω for moments up through second order, which we will use to align a template to the target shape. Given a Borel measurable and Lebesgue integrable function $g : \Omega \rightarrow \mathbb{R}$,

$$\tilde{g}(m, r, U) \triangleq \int_0^1 g(m + t \cdot rU) \cdot \det(\mathbf{I} - t \cdot rS_{\text{rad}}) dt, \quad (3)$$

$$\int_{\Omega} g dV = \int_{\tilde{M}} \tilde{g} dM = \int_{\tilde{M}} \tilde{g} \cdot (U \cdot \mathcal{N}) dA. \quad (4)$$

Here \tilde{M} is the *double* of the medial axis M , indicating that integration is performed over both sides, with \mathcal{N} chosen to point towards the same side as U . The term $dM \triangleq (U \cdot \mathcal{N})dA$ is the *medial measure* defined by Damon [26], which accounts for the failure of U to be orthogonal to M . \mathbf{I} is the identity matrix and S_{rad} is Damon's *radial shape operator*, which measures the rate of change of U along M . In three dimensions S_{rad} can be expressed as a 2×2 matrix computed from derivatives of m and r as described in [24]. Then for simple functions g we can write analytic expressions for \tilde{g} using the *mean radial curvature*, $H_{\text{rad}} \triangleq \frac{1}{2} \text{trace}(S_{\text{rad}})$, and the *Gaussian radial curvature*, $K_{\text{rad}} \triangleq \det(S_{\text{rad}})$:

$$\text{For } g(x) = 1, \tilde{g} = r - r^2 H_{\text{rad}} + \frac{1}{3} r^3 K_{\text{rad}}. \quad (5)$$

$$\begin{aligned} \text{For } g(x) = x, \tilde{g} = m \cdot \left(r - r^2 H_{\text{rad}} + \frac{1}{3} r^3 K_{\text{rad}} \right) \\ + U \cdot \left(\frac{1}{2} r^2 - \frac{2}{3} r^3 H_{\text{rad}} + \frac{1}{4} r^4 K_{\text{rad}} \right). \end{aligned} \quad (6)$$

$$\begin{aligned} \text{For } g(x) = xx^T, \tilde{g} = mm^T \cdot \left(r - r^2 H_{\text{rad}} + \frac{1}{3} r^3 K_{\text{rad}} \right) \\ + (mU^T + Um^T) \cdot \left(\frac{1}{2} r^2 - \frac{2}{3} r^3 H_{\text{rad}} + \frac{1}{4} r^4 K_{\text{rad}} \right) \\ + UU^T \cdot \left(\frac{1}{3} r^3 - \frac{1}{2} r^4 H_{\text{rad}} + \frac{1}{5} r^5 K_{\text{rad}} \right). \end{aligned} \quad (7)$$

The complete volume integrals may now be approximated for any choice of \tilde{g} by

$$\int_{\Omega} g dV \approx \sum_{(m,r,U) \in \mathbb{S}} \tilde{g} \cdot \Delta M, \quad \Delta M \triangleq |U \cdot (m_u \times m_v)| \Delta u \Delta v. \quad (8)$$

These integrals allow the computation of the volume, center of mass, and second order moment tensor of a medially-defined object. Surface integrals over the boundary are even easier. For a Borel measurable function $h : \mathcal{B} \rightarrow \mathbb{R}$,

$$\begin{aligned} \int_{\mathcal{B}} h d\mathcal{B} &= \int_{\bar{M}} \tilde{h} \cdot \det(\mathbf{I} - rS_{\text{rad}}) dM, & (9) \\ &\approx \sum_{(m,r,U) \in \mathbb{S}} \tilde{h} \cdot (1 - 2rH_{\text{rad}} + r^2K_{\text{rad}}) \Delta M, & (10) \end{aligned}$$

where $\tilde{h}(m, r, U) \triangleq h(m + rU)$.

2.2 Single-Subject Model Fitting

Before describing the correspondence optimization, this section outlines the fitting process for a single target shape. First, a similarity transform is applied to the template to align it to the target using its center of mass, volume, and the eigenvectors of the second order moment tensor to define the translation, scale, and rotation, respectively. More robust alignment methods are possible, but this was sufficient for the objects considered in this paper.

We then convert the binary image to a multiscale level-set representation I_{σ} by convolving it with a Gaussian for various choices of σ and choosing a level ℓ_0 for each scale that maximizes the volume overlap with the original image. We deform the template to match the level set at the coarsest scale, and then successively refine it to match each finer scale. Our objective function is given by the squared error integrated over the surface using (10):

$$F_I^{\sigma} = \sum_{(m,r,U) \in \mathbb{S}} (I_{\sigma}(m + rU) - \ell_0)^2 \cdot w_i, \quad w_i = |\det(\mathbf{I} - rS_{\text{rad}}) \Delta M| \quad (11)$$

We use nonlinear conjugate gradient (CG) optimization, with a quadratic penalty function to ensure the model is valid—that is, to ensure that the square root in (1) is real and that no spokes cross inside the object. At points outside the feasible region, the weight w_i is taken to be zero if (1) cannot be evaluated, and the absolute value in w_i handles the case of overfolding. The nonlinear CG method periodically restarts when the next step fails certain orthogonality constraints. We call all the steps between two restarts a *macro step* and hold \mathbb{S} and w_i fixed for its duration. This avoids the need to compute their derivatives and avoids descent towards the trivial global minimum of a zero-volume Ω .

2.3 Correspondence Optimization

In order to optimize the correspondence of a group of medially-defined objects $\mathcal{M}^{(j)}$, we borrow an idea from recent work on computing unbiased, symmetric

atlases [27, 28] and represent the known correspondence via a series of maps ϕ_j that project each shape into a common coordinate system. In the absence of any other prior information, this coordinate system is constructed by Procrustes aligning the target shapes and mapping each point to the linear average of all corresponding points. The map is extended to medially defined points (m, r, U) by projecting $m+rU$ onto the target shape along the U direction, which is normal to the model boundary. This leads to an objective function for correspondence:

$$F_C^j = \sum_{(m,r,U) \in \mathbb{S}} (\phi_j(m, r, U) - \mu_i)^2 \cdot w_i . \quad (12)$$

Here μ_i is the average of the *medially* corresponding points after projection by ϕ_j , that is, those from the same patch, with the same (u, v) coordinates, in multiple models. We use the same sampling \mathbb{S} for each model to ensure this is well-defined, continuing to subdivide in all the models so long as any one violates the threshold in (2). Like w_i , we fix μ_i for the duration of a CG macro step.

In order to avoid interfering with the model fitting process, the deformations allowed in the correspondence optimization are constructed to be approximately transverse to it. That is, we define a *base model* $\mathcal{M}_0^{(j)}$ —taken to be the latest result of the model fitting optimization—and restrict the motion of control points so that their corresponding limit points “slide” along it. For small deformations, within the approximation accuracy of the interpolation, this does not change the shape of the model. The neighborhood of each limit point on $\mathcal{M}_0^{(j)}$ is parameterized using the bilinear map Φ given in [29], which is G^2 and has simple analytic derivatives. Then the control point values that give the desired limit points are recovered by solving a linear system. This system is fixed for a given subdivision surface topology and admits a sparse LU decomposition, whose size is linear in the number of control points [30], making this extremely efficient.

Optimization then proceeds by trading back and forth between model fitting and correspondence optimization. This ensures the approximate transversality holds as the process nears convergence, since this approximation is not very accurate for the initial, large step sizes. First a macro step optimizing F_I^σ is taken for each model, then μ_i is updated and a macro step optimizing F_C^j is taken for each subject j . The two processes do not compete with each other, and so no weight parameter is needed to trade off between the two. The problem is well over-constrained for any modest level of subdivision, so an exact match with the input boundary correspondence will not, in general, be obtained. However, the result conforms to growth or deformation consistent with the medial model, even if the input correspondence does not.

3 Results

We applied our method to a population of synthetic ellipsoids and to a collection of 3D brain objects provided by an ongoing clinical neuroimaging study.

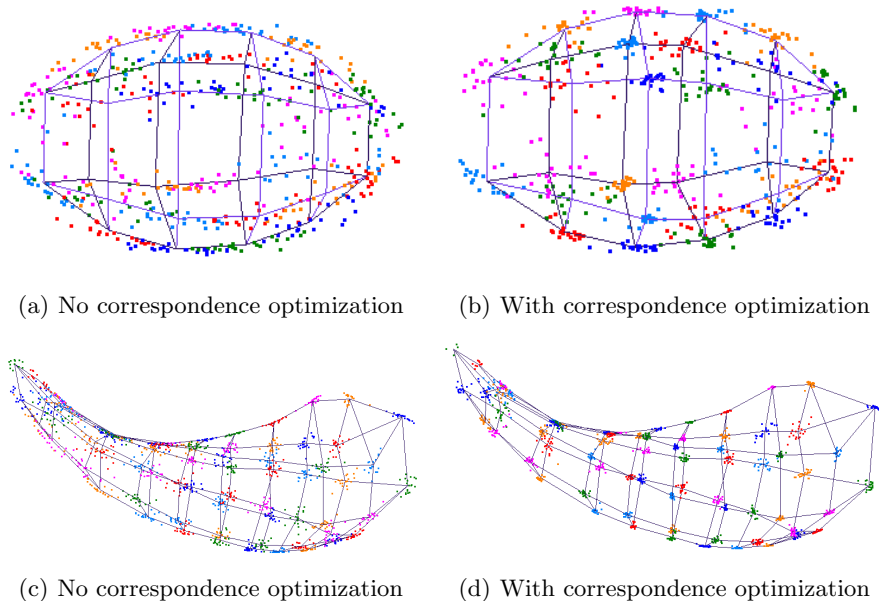


Fig. 3. Clusters of corresponding points in the reference coordinate system, both with and without correspondence optimization enabled. Deformed ellipsoids (top) and left caudate structures (bottom) are shown. Points with the same color belong to the same cluster. The wireframe model connects the cluster centers. Enabling correspondence optimization gives much tighter clusters, especially near the center of the ellipsoid, where the parameterization is most ambiguous. Significant improvement is also shown in the tail region of the caudate.

3.1 Ellipsoid data

A simple test population was created from an ellipsoid deformed by a set of 20 diffeomorphisms of the form [31]

$$\Psi_{\alpha,\beta,\gamma}(x, y, z) \triangleq \begin{bmatrix} x \\ e^{\gamma x}(y \cos(\beta x) - z \sin(\beta x)) \\ e^{\gamma x}(y \sin(\beta x) + z \cos(\beta x)) + \alpha x^2 \end{bmatrix}, \quad (13)$$

where the α , β , and γ parameters control bending, twisting, and tapering, respectively. These parameters were drawn from normal distributions with standard deviations 1.5, 1.05, and 2.12, respectively, and the resulting deformation was applied to a standard ellipsoid with axis lengths of $(1/2, 1/3, 1/4)$ centered around the origin. The result was converted to a $128 \times 128 \times 128$ binary image.

For this data set, ground-truth correspondence is known. Therefore ϕ_j is set to $\Psi_{\alpha_j, \beta_j, \gamma_j}^{-1}$. Although this map does describe the deformation applied to obtain the target object, it may not be physically realistic, since one would not, for example, expect two points on the top and bottom of an ellipsoid with the same x

coordinate to still have the same x coordinate after bending it. We constructed a template by averaging models fit without correspondence optimization, and fit it to the population using 10 scale levels for σ with an iteration limit of $K_{\max} = 200$ macro steps. At each level, the subdivision threshold τ was set to $\frac{1}{2}\sigma$. Each iteration comprised one macro step optimizing the binary image match, F_I^σ , followed by one macro step optimizing the correspondence match for each subject, F_C^j . The optimization only stopped when one of the stopping criteria for CG was encountered for every model, which in this case meant it proceeded to the iteration limit at every scale. The average Dice coefficient of the volume overlap (also computed with a medial integral) was $V_{\text{Dice}} = 96.88\%$, which was actually slightly higher than the 96.72% obtained without correspondence optimization. The continuing evolution of the correspondence optimization even after the image match optimization has converged for a particular subject likely gives it the chance to escape a local minimum and accounts for this small improvement.

In order to visually evaluate how well the surface correspondence was maintained by the medial model, the top row of Fig. 3 shows the endpoints of the spokes associated with each control point mapped into the reference coordinate system via ϕ_j for the models fit both with and without correspondence optimization enabled. When correspondence optimization is enabled, these points form much tighter clusters, especially towards the center of the ellipsoid, where the radius changes more slowly. This is precisely the place where the parameterization is most ambiguous, allowing the correspondence match to produce tight clusters without sacrificing fit quality. The clusters on the ends are not as tight, but for the most part they are still well-separated, unlike their counterparts obtained by fitting models individually.

3.2 Caudate data

Models were also fit to real-world data from an ongoing clinical longitudinal pediatric autism study [32]. Ten subjects each from the autistic and typical groups were chosen and their segmented MRI scans from age four used to test the correspondence optimization. Volumetric segmentations of voxel objects were transformed into surface mesh models and parametrized by SPHARMs [33], with boundary correspondence given by aligning their first-order ellipsoids. This could be replaced by any other correspondence established via analysis of surface geometry or by additional measurements reflecting anatomical or functional geometry. The SPHARM surface was then converted to a triangle mesh and ϕ_j was chosen to map to their average after Procrustes alignment. Ray-triangle intersections computed with the algorithm in [34] were used to project $m + rU$ onto the SPHARM surface, and OBB trees [35] were used to reduced the number of triangles that needed to be tested. To speed up these tests even further, the most recent triangle intersecting each ray was cached, and the OBB tree was searched only if the intersection test against the cached triangle failed. The same scales, iteration limits, and subdivision thresholds were used during optimization as for the ellipse data. The results in the bottom row of Fig. 3 are even more striking than for the ellipse data, especially in the tail region.

4 Conclusion

Sect. 2.1 developed medial integrals as a fundamental tool for taking advantage of the link between the medial axis and the boundary. This allows properties of the boundary to be expressed on the boundary but be evaluated in medial coordinates. This gives continuous medial models advantages of both representations. Although Damon introduced the concept, this work is the first application to demonstrate that they are easily computable and to apply them to computational problems. These medial integrals are applied for computing volume overlap to evaluate goodness-of-fit and for computing second-order moments to align models to a common position, orientation and scale, demonstrating the way medial integrals can be used to compute basic object properties. Most important is the demonstration that these properties can be computed without first converting to a boundary representation, avoiding the additional complexity and approximation error such a conversion would involve.

This paper describes a new correspondence optimization method that works in tandem with the model fitting process to produce a group of models with a common parameterization. Key is a fitting process of a continuous medial model to a population of objects. We show how specific knowledge of correspondence can be incorporated into the model fitting process. This eliminates excess variability in the parameterization of the objects which could mask real statistical effects of the shape change. The correspondence match optimization introduced does not sacrifice the quality of the fit. Instead, it operates transversely to the fitting process up to the tolerance of the model, requiring no tuning parameter to trade off between the two.

The procedure described here serves as a good example of the reasons to use a continuous medial representation over a discrete one and is a non-trivial example of how the link between the medial axis and the boundary can be exploited to give a medial model the advantages of the latter without sacrificing those of the former. The sliding process used to remove ambiguities in the parameterization of the axis would not be possible with a discrete representation. Although the amount of variability this process eliminates is visually remarkable, little has been done to evaluate the quantitative effect this has on the power of statistical tests, and more work needs to be done in this area.

Acknowledgments

This research is supported by the NIH NIBIB grant P01 EB002779. The MRI images of infants, caudate images and expert manual segmentations are funded by NIH RO1 MH61696 and NIMH MH64580.

References

1. Kotcheff, A.C.W., Taylor, C.J.: Automatic construction of eigenshape models by direct optimization. *Medical Image Analysis (MEDIA)* **2**(4) (December 1998) 303–314

2. Davies, R.H., Twining, C.J., Cootes, T.F., Waterton, J.C., Taylor, C.J.: A minimum description length approach to statistical shape modeling. *IEEE Transactions on Medical Imaging* **21**(5) (May 2002) 525–537
3. Heimann, T., Wolf, I., Williams, T., Meinzer, H.P.: 3D active shape models using gradient descent optimization of description length. In: *Proc. of the 19th International Conference on Information Processing in Medical Imaging (IPMI'05)*, Glenwood Springs, Colorado, Springer-Verlag (July 2005) 566–577
4. Cates, J., Meyer, M., Fletcher, P.T., Whitaker, R.T.: Entropy-based particle systems for shape correspondence. In: *Proc. of the International Workshop on Mathematical Foundations of Computational Anatomy (MFCA'06)*, Copenhagen, Denmark (October 2006) 90–99
5. Pizer, S.M., Siddiqi, K., Székely, G., Damon, J.N., Zucker, S.W.: Multiscale medial loci and their properties. *International Journal of Computer Vision* **55**(2–3) (November 2003) 155–179
6. Näf, M., Kübler, O., Kikinis, R., Shenton, M.E., Székely, G.: Shape characterization and recognition of 3D organ shape in medical image analysis using skeletonization. In: *Proc. of the IEEE/SIAM Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA'96)*, San Francisco, California (June 1996) 139–150
7. Zhang, J., Siddiqi, K., Macrini, D., Shokoufandeh, A., Dickinson, S.: Retrieving articulated 3-D models using medial surfaces and their graph spectra. In: *Proc. of the 5th International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, St. Augustine, Florida (November 2005)
8. Malandain, G., Ayache, N., Bertrand, G.: Topological segmentation of discrete surfaces. In: *Proc. of IEEE CVPR'91*, Maui, Hawaii (June 1991) 444–449
9. Golland, P., Grimson, W.E.L., Kikinis, R.: Statistical shape analysis using fixed topology skeletons: Corpus callosum study. In: *Proc. of the 16th International Conference on Information Processing in Medical Imaging (IPMI'99)*, Visegrád, Hungary, Springer-Verlag (June 1999) 382–387
10. Golland, P., Grimson, W.E.L.: Fixed topology skeletons. In: *Proc. of IEEE CVPR'00*. Volume 1., Head Island, South Carolina (June 2000) 10–17
11. Pizer, S.M., Fritsch, D.S., Yushkevich, P.A., Johnson, V.E., Chaney, E.L.: Segmentation, registration, and measurement of shape variation via image object shape. *IEEE Transactions on Medical Imaging* **18** (1999) 851–865
12. Joshi, S.C., Pizer, S.M., Fletcher, P.T., Thall, A., Tracton, G.S.: Multi-scale 3-D deformable model segmentation based on medial description. In: *Proc. of the 17th International Conference on Information Processing in Medical Imaging (IPMI'01)*, Davis, California, Springer-Verlag (June 2001) 64–77
13. Bouix, S., Pruessner, J.C., Collins, D.L., Siddiqi, K.: Hippocampal shape analysis using medial surfaces. *NeuroImage* **25**(4) (May 2005) 1077–1089
14. Styner, M.A., Gerig, G., Joshi, S.C., Pizer, S.M.: Automatic and robust computation of 3D medial models incorporating object variability. *International Journal of Computer Vision* **55**(2–3) (November 2003) 107–122
15. Giblin, P.J., Kimia, B.B.: A formal classification of 3D medial axis points and their local geometry. *IEEE PAMI* **26**(2) (February 2004) 238–251
16. Damon, J.N.: Smoothness and geometry of boundaries associated to skeletal structures I: Sufficient conditions for smoothness. *Annales de l'Institut Fourier* **53**(6) (2003) 1941–1985
17. Damon, J.N.: Smoothness and geometry of boundaries associated to skeletal structures II: Geometry in the Blum case. *Compositio Mathematica* **140**(6) (November 2004) 1657–1674

18. Thall, A.: Deformable Solid Modeling via Medial Sampling and Displacement Subdivision. PhD thesis, University of North Carolina at Chapel Hill (2004)
19. Han, Q., Pizer, S.M., Damon, J.N.: Interpolation in discrete single figure medial objects. In: Proc. of the IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA'06), New York City, New York (June 2006)
20. Yushkevich, P.A.: Statistical shape characterization using the medial representation. PhD thesis, University of North Carolina at Chapel Hill Department of Computer Science (2003)
21. Yushkevich, P.A., Zhang, H., Gee, J.C.: Parametric medial shape representation in 3-D via the Poisson partial differential equation with non-linear boundary conditions. In Christensen, G.E., Sonka, M., eds.: Proc. of the 19th International Conference on Information Processing in Medical Imaging (IPMI'05), Glenwood Springs, Colorado, Springer-Verlag (July 2005) 162–173
22. Yushkevich, P.A., Zhang, H., Gee, J.C.: Continuous medial representation for anatomical structures. *IEEE TMI* **25**(12) (December 2006) 1–18
23. Terriberry, T.B., Gerig, G.: A continuous 3-D medial shape model with branching. In: Proc. of the International Workshop on Mathematical Foundations of Computational Anatomy (MFCA'06), Copenhagen, Denmark (October 2006) 80–89
24. Terriberry, T.B.: Continuous Medial Models in Two-Sample Statistics of Shape. PhD thesis, University of North Carolina at Chapel Hill (2006)
25. Stam, J.: Exact evaluation of Catmull-Clark subdivision surfaces at arbitrary parameter values. In Zorin, D., Schröder, P., eds.: Lecture Notes for the SIGGRAPH'99 Course on Subdivision for Modeling and Animation. (1999) 89–110
26. Damon, J.N.: Global geometry of regions and boundaries via skeletal and medial integrals. Preprint: <http://www.math.unc.edu/Faculty/jndamon/Skel.Str.IV.r.pdf> (2005)
27. Joshi, S.C., Davis, B., Jomier, M., Gerig, G.: Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage; Supplemental issue on Mathematics in Brain Imaging* **23**(Supplement 1) (April 2004) S151–S160
28. Škrinjar, O., Tagare, H.: Symmetric, transitive, geometric deformation and intensity variation invariant nonrigid image registration. In: Proc. of IEEE ISBI'04, Arlington, Virginia (April 2004) 920–923
29. Loop, C.: Second order smoothness over extraordinary vertices. In: Proc. of the 2004 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing, Nice, France (July 2004) 165–174
30. Halstead, M., Kass, M., DeRose, T.: Efficient, fair interpolation using Catmull-Clark surfaces. In: Proc. of the 20th Annual SIGGRAPH Conference on Computer Graphics and Interactive Techniques, Anaheim, California (August 1993) 35–44
31. Han, Q. Personal communication (2006)
32. Styner, M.A., Gorczowski, K., Fletcher, P.T., Jeong, J.Y., Pizer, S.M., Gerig, G.: Statistics of pose and shape in multi-object complexes using principal geodesic analysis. In: Proc. of the Third International Workshop on Medical Imaging and Augmented Reality (MIAR'06), Shanghai, China, Springer-Verlag (August 2006) 1–8
33. Brechbühler, C., Gerig, G., Kübler, O.: Parameterization of closed surfaces for 3-D shape description. *Computer Vision, Graphics, and Image Processing: Image Understanding* **61**(2) (March 1995) 195–170
34. Möller, T., Trumbore, B.: Fast, minimum storage ray/triangle intersection. *Journal of Graphics Tools* **2**(1) (October 1997) 21–28
35. Gottschalk, S., Lin, M.C., Manocha, D.: OBB-tree: A hierarchical structure for rapid interference detection. In: Proc. of the 23rd Annual SIGGRAPH Conference on Computer Graphics and Interactive Techniques. (August 1996) 171–180