

Nested Sphere Statistics of Skeletal Models

Stephen M. Pizer, Sungkyu Jung, Dibyendusekhar Goswami, Jared Vicory, Xiaojie Zhao, Ritwik Chaudhuri, James N. Damon, Stephan Huckemann and J.S. Marron

Abstract We seek a form of object model that exactly and completely captures the interior of most non-branching anatomic objects and simultaneously is well suited for probabilistic analysis on populations of such objects. We show that certain nearly medial, skeletal models satisfy these requirements. These models are first mathematically defined in continuous 3-space, and then discrete representations formed by a tuple of spoke vectors are derived. We describe means of fitting these skeletal models into manual or automatic segmentations of objects in a way stable enough to support statistical analysis, and we sketch means of modifying these fits to provide good correspondences of spoke vectors across a training population of objects. Understanding will be developed that these discrete skeletal models live in an abstract space made of a Cartesian product of a Euclidean space and a collection of spherical spaces. Based on this understanding and the way objects change under various rigid and nonrigid transformations, a method analogous to principal component analysis called *composite principal nested spheres* will be seen to apply to learning a more efficient collection of modes of object variation about a new and more representative mean object than those provided by other representations and other statistical analysis methods. The methods are illustrated by application to hippocampi.

Stephen M. Pizer, Dibyendusekhar Goswami, Jared Vicory, Xiaojie Zhao, Ritwik Chaudhuri, James N. Damon and J.S. Marron,
University of North Carolina at Chapel Hill, NC, USA. e-mail: pizer@cs.unc.edu

Sungkyu Jung
University of Pittsburgh, PA, USA. e-mail: sungkyu@pitt.edu

Stephan Huckemann
Georg-August-Universität, Göttingen, Germany. e-mail: huckeman@math.uni-goettingen.de

1 Object models suitable for statistics

The objective of statistics on a population of 3D (or 2D) objects is to produce such entities as the mean object and a shape space spanned by principal modes of variation of the object. These statistical descriptors are derived from training cases, each typically given by a pile of object boundary contours. Statistics on objects has been applied to quite a variety of types of geometric model derived from the boundary cases:

1. Boundary point distribution models ([4], [8], [17]). This popular form of model has been analyzed by principal component analysis (PCA), although Kendall [16] has shown this is not strictly appropriate because such models live in a space formed by a Cartesian product of a Euclidean space and a quotient of a high-dimensional sphere modulo the rotation group. However, even the correct analysis ignores the shape of the object interior and, partially as a result, has difficulty limiting its shape space to models whose boundaries do not cross themselves.
2. Deformation-of-atlas models, wherein the displacement of each voxel in the atlas is provided [21], [1]). These models have an enormous dimension, with the result that the statistical analysis is expensive and unstable with respect to the sampling into training cases and to noise in those training cases.
3. Implicit models, such as level surfaces of pseudo-signed-distance functions [18], [31] that do not live in Euclidean feature spaces but are often analyzed (by PCA) as if they did.
4. Skeletal models. While skeletal modeling is common, as summarized in the book by Siddiqi and Pizer [27], work on statistics of skeletal models has been the major topic of only one group, namely ours. That work emphasized that skeletal models, like those in the first two categories, live on abstract manifolds that are curved. Work on skeletal and other object models that live on abstract manifolds have been the subject of study of Fletcher [10]. Our previous work on skeletal model statistics, on which Fletcher was an early contributor, is laid out in chapters 8 and 9 of Siddiqi and Pizer [27]. The strength of skeletal models is that they richly model the object interior and boundary and that they yield an object-relative coordinate system for the object interiors. The difficulty with general skeletal models, and especially the medial models that are their most common form, is that even for objects that do not intuitively feel like a tree of figures, the model takes the form of the tree. However, the branching structure of the tree does not stay the same over the object population. This makes statistical analysis of these models very difficult. In our previous work we have solved this problem by using skeletal models whose boundaries are only approximately that of the training cases but which have no branches.

From the previous discussion one can conclude that no form of model has been introduced that accurately and richly represents the training data and truly is fully appropriate for statistical modeling. In this paper, we define a form of skeletal model that precisely captures most non-branching anatomic objects and a form of statistical analysis more well suited to these models than ones that have been previously ap-

plied. The result of this statistical analysis, we show by illustration, produces more appropriate object means, lower-dimensional shape spaces, and more descriptive modes of variation than PCA-based statistical analysis of other skeletal or nonskeletal models. While our object modeling approach allows us to capture branching objects as well, the statistical techniques available to us or others is not yet capable of handling structures with variable branching structure. Therefore, here we do not attempt here to model all object populations.

2 Skeletal models of non-branching objects

Skeletal models capture the interior of objects (the egg), and as such they more stably and richly capture object shape than models that capture only the boundary of objects (the eggshell). We define a continuous interior-filling skeletal model of an object as a locus of spoke vectors $(\underline{p}, \mathbf{S})$ with tail at \underline{p} and tip at $\underline{p} + \mathbf{S}$ such that

1. none of the vectors cross each other;
2. the union of the vectors forms the interior of the object and the union of the spoke tips forms the boundary of the object. The consequence of conditions 1 and 2 is that each point in the object is reached by precisely 1 spoke.
3. the union of the tails, which we call the “skeletal locus”, forms a fully folded, multi-sided locus (double-sided for slab-shaped objects). That is, except at the fold, in slabs each position p appears twice in the set of spokes (see Fig. 1). The reader should not confuse this with the more traditional characterization of the skeletal locus as an unfolded point set that happens to have two spokes at every point but the end points.

We call such a representation an interior-filling s -rep, or for this paper just an s -rep. We will refer to the whole s -rep by the notation \underline{m} . In this paper we restrict ourselves to s -reps of 3D objects that have the topology of a sphere (have no interior cavities or through holes).

We define a non-branching 3D object as a region of space for which there exists an s -rep with a single cyclic fold curve, i.e., for which the skeleton does not branch. Examples of such objects are the kidney and the esophagus (into which we swallow).

Medial models, as defined by Blum [2] are a form of skeletal model, but they cannot exactly capture the entire interior of most non-branching anatomic objects without an arbitrary number of skeletal branches and variation of the number of branches and their branching structure. These render the




Fig. 1 An interior-filling skeletal model in 2D. The two sides of the skeletal locus are shown slightly separated, but in fact they are coincident. In the continuous form there is a spoke vector at every point on the folded curve forming the skeleton.

model unsuitable for probabilistic analysis. Instead we seek a skeletal model [6] that allows overcoming this branching instability. In particular, our target is that all instances of the object have the simplest, stable branching pattern, i.e., without branches. In these s-reps the skeletal locus is *smooth* except at the fold points.

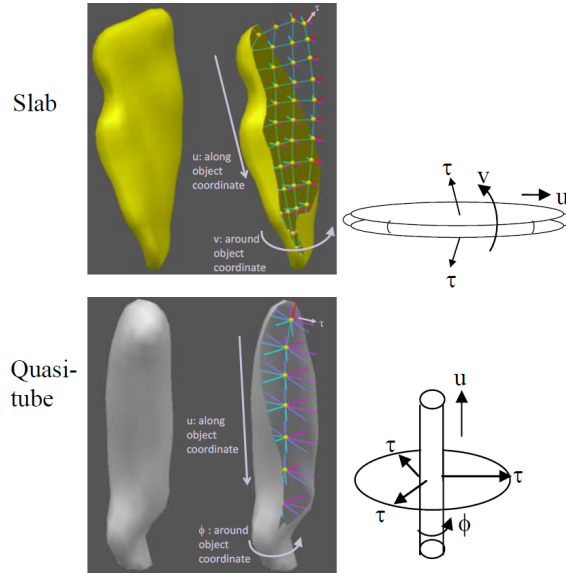


Fig. 2 Slabular and quasi-tubular s-reps and their parameter spaces. In the slab, $\underline{u} = (u, v)$ parametrize by mapping the long (u) and middle (v) dimensions of the object to the respective dimensions of an ellipsoid collapsed along its short axis, which in turn is mapped to the sphere by mapping its crest to the equator, its north side to the sphere's northern hemisphere, and its south side to the sphere's southern hemisphere. In the quasi-tube, $\underline{u} = (u, \phi)$ parametrize by mapping the along-tube dimension (u) to the axis of a cylinder with infinitesimal radius and mapping the around-tube dimension (ϕ) to the normal directions of the infinitesimal cylinder. In both forms of s-reps, τ parametrizes fractional distance along the spokes, where $\tau = 0$ at the skeletal locus and $\tau = 1$ at the object boundary. In both forms (u, τ) parametrize the interior of the object.

Such s-reps capture all non-branching anatomic objects in the human body. This includes not only slabular objects such as the kidney, most muscles, the heart, and the bladder, and quasi-tubular objects, such as the esophagus and the rectum, but also folded slabular objects such as the cerebral cortex and folded quasi-tubular objects such as the intestines. These models capture objects with small pimples and dimples; the main objects not captured are objects with hook-shaped or curvy protrusions and indentations and those that branch multiple times.

As illustrated in Fig. 2, there are two types of s-reps. In the first type, which we call *slabular*, the smooth points form two copies of a smooth sheet, with the two sheets pasted to each other, and the fold points form themselves into a single smooth cyclic curve, called the *end curve*. In the second type, which we call *quasi-tubular*, the smooth points take the form of a bent cylinder of infinitesimal radius, i.e., a cylinder whose surface is made up of copies of an axis curve with all the copy curves pasted to each other, and the fold points consist of two discrete points, called the *end points*. In both types a continuum of spokes, formed by vectors with tails on the skeletal locus and tips on the object boundary, cover the interior of the object.

In slabular s-reps the skeletal locus forms a two-sided sheet $\underline{p}(\underline{u})$ with a cyclic fold curve. The parametrization \underline{u} has the topology of a sphere. In this work, for slabular s-reps we will parametrize by the sphere and let the equator of the parametrizing sphere map onto the fold of skeletal locus. Thus we will refer to the north and south sides of the sheet. For each \underline{u} a spoke $\mathbf{S}(\underline{u})$ exists with its tail on $\underline{p}(\underline{u})$ and its tip on an object boundary point $\underline{b}(\underline{u})$. \underline{u} parametrizes the skeletal locus along the object and across and around the object; it also parametrizes the spokes $\mathbf{S}(\underline{u})$, which pass within the object, as well as the object boundary $\underline{b}(\underline{u})$. The lengths of the spokes, $|\mathbf{S}(\underline{u})|$, which we will call $r(\underline{u})$, and the directions of the spokes $\mathbf{U}(\underline{u}) = \mathbf{S}(\underline{u})/r(\underline{u})$ are also defined. Thus \underline{u} parametrizes the whole s-rep $\underline{\mathbf{m}}(\underline{u}) = (\underline{p}(\underline{u}), \mathbf{U}(\underline{u}), r(\underline{u}))$.

In quasi-tubular s-reps the skeletal locus forms a collapsed (infinitesimal radius) bent cylinder with hemispherical caps at both ends, such that the axis of the cylinder is a space curve and the spokes emanating from each axis-orthogonal cut of the cylinder end in a common plane. Here the parametrization \underline{u} is typically on the unit-radius, possibly hemispherically capped right-circular cylinder (Fig. 2); it parametrizes the axis of (forming) the skeletal locus by the along-cylinder scalar variable u , and it parametrizes the angle around the circular tubular cross sections by another cyclic scalar variable ϕ . As with the slabular s-rep, \underline{u} parametrizes the spokes $\mathbf{S}(\underline{u})$, their directions $\mathbf{U}(\underline{u})$, their lengths $r(\underline{u})$, and the object boundary $\underline{b}(\underline{u})$. Quasi-tubular s-reps are used to model roughly tubular objects such as the esophagus, but they can also be used to model objects like muscles whose cross-sections are star-shaped but not roughly tubular.

For both slabular and quasi-tubular s-reps the position along a spoke from its tip to its tail is parametrized by the proportion-of-spoke-length variable τ , so the union of the interior and the boundary of the object is parametrized by (\underline{u}, τ) . In the remainder of this paper we restrict ourselves to slabular s-reps.

The mathematical theory of skeletal models is presented in Damon's chapter 3 of the book *Medial Representations* by Siddiqi and Pizer [27] and more completely in Damon's papers: [7, 6]. There the geometry of our class of s-reps is shown to be neatly divisible into a) the differential geometry of the two sides and fold of the skeletal locus $\underline{p}(\underline{u})$ and b) a variant of differential geometry, which we call radial geometry, of the (tail-less) spokes $\mathbf{S}(\underline{u})$ with respect to motions on the tangent plane of the skeletal locus. The radial geometry of the spoke directions $\mathbf{U}(\underline{u})$ is of special interest; it can be divided into a *radial shape operator* $S_{rad}(\underline{u})$ applicable at non-fold medial points and an *edge shape operator* $S_E(\underline{u})$ [for \underline{u} on the equator of the parametrizing sphere or the ends of the parametrizing cylinder] applicable at the fold curve. Each of these operators are represented by a 2×2 matrix. These shape operators describe $\mathbf{U}(\underline{u})$ motion in a way analogous to the way the ordinary shape operator from differential geometry describes the motion of boundary normals. There are special matters dealing with the appropriate coordinate system in which to describe $\mathbf{U}(\underline{u})$ motion and the need for projection of that motion onto the tangent plane to $\underline{p}(\underline{u})$, but we will not go into those in this paper.

Damon shows that an interior-filling locus $(\underline{p}(\underline{u}), \mathbf{U}(\underline{u}), r(\underline{u}))$ has no locally crossing spokes (is *geometrically legal*) if and only if for all \underline{u} the positive eigenvalues of $S_{rad}(\underline{u}) < 1/r(\underline{u})$ and any positive generalized eigenvalue of $(S_E(\underline{u}), J) < 1/r(\underline{u})$,

where J is the 2×2 matrix all of whose elements are 0 but for the upper left element, which is 1.

Discrete s-reps sample $\mathbf{m}(\underline{u})$ into a network of spoke vectors. In our work with slabular s-reps the sampling is into a north $m \times n$ grid of spokes, a south $m \times n$ grid of spokes, for some m, n , and a cyclic chain of $2m + 2n - 4$ edge (crest) spokes (Fig. 2). These are organized into an $m \times n$ grid of “skeletal atoms”, each consisting of 2 or 3 spokes with a common tail position. For interior grid points there are 2 spokes, one on the north side and one on the south side. For exterior grid points there are 3 spokes, one on the north side, one on the south side, and a crest spoke bisecting the other two.

Interpolation of a continuous s-rep from a discrete s-rep proceeds according to the method of Han [11], in which $\underline{p}(\underline{u})$, $r(\underline{u})$, and $\mathbf{U}(\underline{u})$ are interpolated in an intertwined way. The spokes are interpolated using the fact that $S_{rad}(\underline{u})$ or $S_E(\underline{u})$ times a small step vector on the tangent plane to the skeletal locus at $\underline{p}(\underline{u})$ allows the calculation of the spoke direction swing corresponding to that step. The $rS_{rad}(\underline{u})$ or $rS_E(\underline{u})$ matrix is first interpolated and then the implied spoke swing and length scaling for various small steps is integrated. The skeletal axis positions $\underline{p}(\underline{u})$ are interpolated using a Hermite-like interpolation using both the discrete positions and skeletal sheet normals there, where the normal is computed as the difference of the two spoke direction vectors with a common skeletal sheet position. This relation for the normal holds only for Blum medial sheets, but given the “as medial as possible” criterion for skeletal sheets discussed below, it is sensible to use it for the skeletal sheet interpolation.

3 Obtaining s-reps suitable for probabilistic analysis

There are three properties necessary for making a population of s-reps suitable for probabilistic analysis. The first is that the branching topology of all members of the population should be the same. There is nice, albeit not yet usable, work on statistics of trees with variable branching topology [28, 30], but this is beyond the scope of this paper. In this paper we restrict the branching situation further to having no skeletal branches at all. This requires branching objects such as blood vessels to be treated by compositing single branch components, but this, too, is beyond the scope of this paper. This is the reason for the restriction mentioned in section 1 to skeletal loci that are smooth except at the fold curve.

The second property needed to do probabilistic analysis is that small deformations in the objects yield small deformations in the skeletal model. This can be made to occur when the process of fitting s-reps into the data describing the object has certain behaviors. These will be discussed in section 3.1.

The third property needed to do probabilistic analysis is that each spoke in one s-rep in the training population be at a position, orientation, and length that geometrically corresponds to that of the same spoke in the other training s-reps. Accomplishing this is discussed in section 3.2.

3.1 Fitting unbranched s-reps to object description data

Typically object data is provided in the form of a tiled object boundary, a pile of parallel planes each containing a boundary contour, or a binary discrete image. To fit s-reps to such data, we first transform the data into a signed distance image, where the distance is negative inside the object and positive outside and where the zero level curve of distance represents the object boundary. We have methods, using constrained flow controlled by the Laplacian of curvature, that yield a distance image representing a smooth boundary staying within a half-voxel of the object data [22].

Our objective is to study shape variability by statistical means, so the pair of distance images for many population samples of our object will be provided. For any object sample there are very many s-reps that can be fitted to this data. In particular, the skeletal locus could in principle be anywhere across the short dimension of the object. In the extreme it could be along one side of the object, such that the object was covered by lines going from one side of the object (e.g., the south side) to the other side of the object. Since our goal is to do statistical analysis of the fitted s-reps, we need to fit s-reps that have as little variability over the training cases as possible.

The closer the skeletal locus is to bisecting the top and bottom sides, the wider the range of pimples and dimples on the boundary that could be accommodated. More importantly, we can more stably constrain the geometry of the skeletal locus if we make it bisect the two sides. However, as indicated earlier, the Blum medial locus will not do, because it generates branches for pimples and dimples. So in brief, we want a skeletal locus smooth but at the folds that is space filling and, within those constraints, *as medial as possible*. The following section gives properties of truly medial fits that can be used in measuring how medial an s-rep is. The properties are described in the form of penalties that contribute to a weighted sum that is minimized in the fitting process.

3.1.1 As medial as possible

The following medial properties should be approximate targets for an as-medial-as-possible s-rep. Each property yields a penalty term in the aforementioned weighted sum. Some depend on the input signed distance image $D(\underline{x})$, and some involve the relation between the north (south) side spoke $\mathbf{S}(\underline{u})$ and $\mathbf{S}(\underline{v})$, its counterpart on the other side of the (double-sided) skeletal sheet with the same $p(\underline{u})$ value.

1. Equal $r(\underline{u})$ values for the northside and southside spokes at the same $p(\underline{u})$: in the fitting of a discrete s-rep we penalize according to $|r(\underline{u}) - r(\underline{v})|$.
2. Spokes are boundary normals: in the fitting we penalize according to the deviation from boundary normality, i.e., $\cos^{-1}(\widehat{\nabla}D(\underline{x}) \cdot \mathbf{U}(\underline{u}))$ at the spoke end, where, given a vector \mathbf{W} , $\widehat{\mathbf{W}}$ is the unit vector in the direction of \mathbf{W} .
3. The (interpolated) skeletal locus $p(\underline{u})$ bisects $\mathbf{S}(\underline{u})$ and $\mathbf{S}(\underline{v})$. Equivalently, the normal to $p(\underline{u})$ is in the same direction as $\mathbf{U}(\underline{u}) - \mathbf{U}(\underline{v})$: in the fitting, for each spoke we penalize according to the angle between the normal to $p(\underline{u})$ and

$\mathbf{U}(\underline{u}) - \mathbf{U}(\underline{v})$. This penalty has not yet been implemented in our code to fit s-reps to objects.

4. End curves' spokes are at crests of the zero level curve of the distance image that forms the object boundary; this involves three properties, all involving the principal directions w_1 and w_2 and the associated principal curvature κ_1 of the level surface at the spoke end, where κ_1 is the lesser (more negative, i.e., more convex) principal curvature:
 - a. The spoke end is at an extremum of κ_1 along the integral curve of w_1 , so the directional derivative of κ_1 in the w_1 direction should be zero. In the fitting we have attempted to penalize according to the magnitude of that directional derivative. However, the high order derivative of the distance function that this involves proved to be unstable, so for now this term is not included in the penalty list.
 - b. The plane of the spoke and its infinitesimally near neighbors along the principal curve of w_1 should be orthogonal to w_2 . We penalize according to the angle between w_2 and $\mathbf{U}(\underline{u}) \times \mathbf{U}(\underline{v})$, the normal to the plane of the spokes at that end curve point.
 - c. The boundary implied by the s-rep at the spoke end and the level surface of the distance at the spoke end should agree in the first principal curvature. We have not yet implemented a penalty for deviating from this condition, but in a future implementation we plan also to penalize according to the magnitude of the difference between the two relevant radii of curvature.




Fig. 3 A 2D object with multiple vertices (the 2D correspondent to a crest in 3D), together with the desired skeletal sheet.

This method of forcing spokes into the distance function's crest is good when there is a strong crest for the end curve to fit into, but in some objects this crest is weak or, even worse, the object has more than one crest and we wish the s-rep to fit between them. This situation is illustrated in Fig. 3, where the skeletal sheet bisects the twin crests on the right side of the object, instead of favoring one crest over the other.

3.1.2 Fitting s-reps to distance images

Fitting s-reps to such signed distance images is done in five stages.

In the first stage, to align the training cases, a reference s-rep is translated, rotated, and possibly uniformly scaled into the space of the distance image via matching of moments of boundary positions or via minimizing sum of squared distances between designated s-rep spoke ends and object landmarks provided by the user.

In the second stage we restrict the corresponding northside and southside spokes to having the same length and move each atom into place while maintaining a reg-

ular array that is geometrically legal. In the third (spoke) stage we adjust the spoke lengths and angles to match the object boundary quite closely. After the PCA-like statistical analysis described in section 5 on the results of the third stage, the fourth stage obtains better correspondence among the fits to the various training cases. It accomplishes that by restricting the fits to the shape space resulting from the statistical analysis. In the fifth stage we tighten the fit to the distance function by allowing small deviations from the shape space; in particular, we again adjust the spoke lengths and angles to match the object boundary quite closely. The following describes these stages in more detail.

In the second stage, iteratively over the skeletal atoms (each consisting of a spoke duple or triple with a common hub) we optimize an objective function [19] that is a sum of two major penalty terms summed over the 2 or 3 spokes. Each major term is a sum of subsidiary terms. The first major term measures how well the candidate spokes are fit into the distance image; the second of these major terms measures the geometric fitness of the candidate s-rep. Both terms include penalties for deviation from being medial.

For the spoke $\mathbf{S}(\underline{u})$ the spoke-to-object-mismatch penalty sums terms penalizing misfit to $D(\underline{b}(\underline{u}))$ (the signed distance value at the spoke end) and its derivatives. The geometric fitness penalty sums terms penalizing general geometric misbehavior such as lack of adherence to the medial properties.

Spoke-to-object mismatch penalties: All of these penalties are summed over spokes interpolated from the original discrete set. In our application to hippocampi, each quad of 4 spokes is subdivided by 4 in each dimension, producing a quad of 25 interpolated spokes. No interpolation is done, at present, around the crest.

1. Zeroth order fit to the distance image: $|D(\underline{b}(\underline{u}))|$.
2. First order fit to the distance image: the angle between $\mathbf{S}(\underline{u})$ and $\nabla D(\underline{b}(\underline{u}))$, as described in section 3.1.1, item 2.
3. The crest property penalties described in section 3.1.1, item 4.

Geometric fitness penalties:

1. S-rep irregularity: To achieve roughly uniform coverage of the object interior and, given the crest-fitting of the end spokes, to provide approximate correspondence of spokes across cases, a penalty for each spoke is applied proportional to its difference from the average of its neighbors. At present, the average is calculated in a way consistent with the principal geodesic analysis approach described in section 5.1.1.
2. Illegality of spoke crossings: a penalty is added for S_{rad} or S_E values for a spoke being greater than the reciprocal of their spoke lengths (see section 2).
3. Deviation from medial properties, as described in items 1 and 3 of section 3.1.1.
4. The distance between $\mathbf{S}(\underline{u})$ and the result of the first stage for that spoke. The distance measure reflects positional deviations at the spoke end.

In the third, i.e., spoke, stage the spoke tail $p(\underline{u})$ is fixed, and the optimization is only over the spoke direction $\mathbf{U}(\underline{u})$ and the spoke length $r(\underline{u})$. The spoke-to-object

mismatch penalties are as above, but the geometric fitness penalty is only on the magnitude of the difference of both the spoke direction and the spoke length from the result at the first stage.

In the fourth stage the optimization is over the shape space determined by the statistical analysis of those third stage fits that are adequately good, and the initialization is always from the mean s-rep of those fits. In that stage the geometric atypicality penalty is the negative log prior density (up to a constant multiplier and additive constant) of being in a population of s-reps, given by the Mahalanobis distance function also determined from the statistical analysis. In sections 5.2 and 5.3 we describe how to compute a set of eigenmodes \mathbf{v}^i of s-reps and associated principal variances σ_i^2 , using the actual abstract space in which an s-rep falls. Any s-rep in that population is designated by a tuple of scores $\underline{\alpha}$ of the respective eigenmodes, and $\sum_i (\alpha_i^2 / \sigma_i^2)$ forms the Mahalanobis distance of that s-rep.

The fifth, spoke, stage is just like the third stage, but on and relative to the results of the fourth stage.

In all of the stages conjugate gradient optimization of the objective function is applied.

Fig. 4 shows some sample hippocampus fits by this method, as described in section 6.




Fig. 4 S-reps fitted to 9 of the 62 hippocampi used for the statistical analysis (see section 6)

3.2 Achieving correspondence of spoke vectors

The previous work on achieving correspondence in training populations of anatomic object models has involved shifting points on object boundaries in point-distribution models (PDMs) [3, 20] [5]. We find the most attractive method is that of [3, 20], which minimizes a sum of two entropy terms: $H(\underline{\mathbf{z}}) - \alpha \sum_i H(\underline{\mathbf{x}}_i)$. The first term measures the tightness of the probability distribution on the representation entities, here possibly including not only the boundary points but also derived values such as curvature. The second term sums negatives of entropies over the training cases, with each entropy measuring the uniformity of the distribution of the points on the respective object boundary. We are in the process of extending this method to spoke correspondences. In our case the representation entities are spokes, and the entropy possibly includes not only skeletal sheet curvatures (eigenvalues of the shape operator at spoke tails) but also spoke direction curvatures (eigenvalues of the S_{rad} or S_E operators of spokes). Also, the uniformity measurement must be of spokes within the object volume.

However, this method has not been fully implemented at the time this paper is written. Thus the following statistics are based on fits whose correspondence come from common shape-based penalties used for fitting the training cases and the common shape space in which they fall before refinement by the final spoke stage.

4 The abstract space of s-reps and common configurations of s-rep families in that space

4.1 The abstract space of s-reps

Let each s-rep $\underline{\mathbf{m}}^k$ in a training population consist of n spoke vectors $\{(p_i, r_i, \mathbf{U}_i) \mid i = 1, 2, \dots, n\}$ that correspond across the population. The set $\{p_i\}$ of points on each training discrete s-rep's skeletal locus form a PDM (point distribution model) that can be centered so that its center of gravity is the origin. Also, the result can be scaled by a factor making the sum of the squared distances of the points to the origins equal to unity. Thereby each PDM is describable by a scaling term represented by the log of its scale factor γ and a tuple of scaled, centered spoke tail points that abstractly live on the unit $3n - 4$ dimensional sphere S^{3n-4} .

Each spoke direction \mathbf{U}_i abstractly lives on the unit 2-sphere. Each $\log r_i$ value, as well as $\log \gamma$, abstractly live in Euclidean space. Thus a discrete s-rep lives in $\mathcal{R}^{n+1} \times S^{3n-4} \times (S^2)^n$, i.e., the Cartesian product of $n + 2$ manifolds, one of which is Euclidean and all of the rest of which are spheres. If the number of training cases N is less than $3n - 4$, the space of the scaled, centered spoke tails is restricted to the intersection of S^{3n-4} with the Euclidean space of dimension $N - 1$ passing through the center of S^{3n-4} , i.e., in S^{N-1} . Also in that case, the n radii live in \mathcal{R}^{N-1} . In the

experiment on hippocampi reported in section 5, $N = 62$ and $n = 66$. Thus each hippocampus s-rep is a point in the composite space $\mathcal{R}^{62} \times S^{61} \times (S^2)^{66}$.

4.2 Families of s-rep components on their spheres

The statistical analysis method we will describe in section 5 involves includes analysis on each of the here, 67 component spherical spaces, then compositing the analyzed information with one another and with the Euclidean space data, and analyzing the composited data. Therefore, it is of interest to consider how the s-rep data points distribute themselves on each sphere.

Certain basic transformations of the objects can be expected: global rigid rotations, rotational folding about an axis, and twisting about an axis. We have shown [24] the following behaviors of the spoke-direction points on S^2 under these transformations. For each transformation each spoke moves on a small circle on S^2 about the rotational axis of the transformation; different spokes move on different small circles, but all the circles share the same axis. In rotational folding and twisting different spoke-direction points rotate in opposite directions as the fold or twist progresses.

Also, for rigid rotation the points on S^{3n-4} or S^{N-1} (here S^{61}), each describing the tuple of skeletal points for a separate training s-rep, move on a small circle (1-sphere).

We have experimentally confirmed (see Fig. 5) the relevance of these ideas, by observing that real data s-rep spokes tend to arrange themselves near small circles.

We will now examine methods for statistically analyzing such data on a d -dimensional sphere in a PCA-like fashion, to describe the data by a suitable notion of mean and a shape space formed from a limited number of eigenmodes of variation.




Fig. 5 Data points for an s-rep spoke in 62 hippocampi (see section 6)

5 Training probability distributions in populations of discrete s-reps

5.1 Previous methods for analyzing data on a sphere

To understand the methods we will describe, it is useful to describe two approaches to PCA in a d -dimensional Euclidean space. The two approaches are called “forward” and “backward” according to the order in which the dimensions are analyzed. In a Euclidean space the two approaches are equivalent, but on spheres (and indeed any curved space) they can give far from the same results.

In the forward approach, the mean is computed as the 0-dimensional space (point) best fitting the d -dimensional data. Then the first principal component is computed as the 1-dimensional space (line) best fitting the d -dimensional data. Then the second principal component is computed by first finding the 2-dimensional space (plane) best fitting the d -dimensional data and then taking the direction in that plane orthogonal to the first component. And so on. The analysis proceeds from dimension 0 to 1 to 2 to ... to $d - 1$. At each step the fit is to the d -dimensional data.

The backward approach begins by fitting the best hyperplane of dimension $d - 1$ to the d -dimensional data. *Then the data is geodesically (orthogonally) projected onto this hyperplane.* Next the fit of a hyperplane of dimension $d - 2$ is made to this projected data. And then the data from the hyperplane of dimension $d - 1$ is projected onto this hyperplane of dimension $d - 2$. And so on, until we have a line fitted to the data that has been projected on a plane and we project the data from the plane onto the line. Finally, the mean is the 0-dimensional best fit to the data on the line. In this method the fit at the k th step is to the data projected onto a space of dimension $d - k + 1$.

In PCA each training point can be expressed as the mean plus a weighted sum of the d eigenmodes. The weights are called scores. The score of eigenmode $d - k + 1$, $k = 1, 2, \dots, d$ for that training point is the signed projection distance of that data point from the space of dimension $d - k + 1$ to the space of dimension $d - k$.

5.1.1 Principal Geodesic Analysis

The method we had previously used to analyze m-reps (s-reps with equal-length spokes at common spoke tails) was the Principal Geodesic Analysis (PGA) method of Fletcher [9]. This method is a forward approach. It begins with computing the intrinsic forward Fréchet mean on each component, spherical and Euclidean, in the composite space. The mean on a sphere is the point on the sphere whose sum of squared geodesic distances to the data points is minimum. The intrinsic Fréchet mean is a forward mean because the best zero-dimensional approximation to the data, the mean, is computed relative to the full d -dimensional sphere.




Fig. 6 Data clustered on a sphere

Given this mean, the form of PGA that we use moves the data from its manifold onto a tangent hyperplane at the mean with a Euclidean metric determined by keeping directions and the condition that every point on the sphere has the same Euclidean distance to the mean on the tangent hyperplane as its geodesic spherical distance to the mean on the sphere. On the data on this tangent plane the method carries out PCA, and its eigenvectors are mapped back to the original sphere, yielding geodesics through the mean. Strictly speaking, the method is applied to the composite space, not sphere by sphere, but it could be applied sphere by sphere.

This method works very well when the data cluster on each of the composite spheres (and the Euclidean space in the composite) because then data projected on

the tangent plane represents the original data well (Fig. 6). However, in our case the data frequently live on small circles that are far from the pole, whereas the mean may be found at the pole. It is then not surprising that PGA will not be optimal.

5.1.2 GPCA: Geodesics with the Mean on the First Geodesic

Huckemann et al. [12, 13] proposed a method called geodesic principal component analysis (GPCA) which inspired our final method by having an important backward component. They realized that the Fréchet mean may be far from the best fitting great circle. Hence, if the data lives on or near a circle on a sphere, the mean should be computed on that circle after projection of the data onto it. Their method found the best fitting geodesic (great) circle on the sphere, projected the data onto that, and found the Fréchet mean on that great circle of the projected points. This is a backward mean, since it computed after projection from the original d -space onto a 1-dimensional subspace.

The method of Huckemann et al. on a d -sphere then went on to compute principal geodesic circles through the mean in a forward fashion, in a way quite similar to PCA, albeit on a sphere. Its weakness for our data comes from two problems. First, it fits great circles, whereas our data live near small circles. Second, it is backward only at the first step, whereas the fully backward method of Jung [14] described next appears superior.

5.1.3 Composite Principal Nested Spheres

A method better recognizing the loci along which the s -reps tend naturally to vary is the Composite Principal Nested Spheres (CPNS) method of [14]. It consists of 1) component by component estimation of principal modes of each of the spherical manifolds separately, each producing scores that are suitable for Euclidean analysis; then 2) compositing those analyses with one another and the Euclidean data, followed by PCA on the result.

5.2 *Training probability distributions on s -rep components living on spheres: Principal Nested Spheres*

We begin by describing the principal nested spheres (PNS) approach that is applied to the data on each sphere separately. This method has already been described in [14]. As suggested, the approach is backwards analysis, and at each step the best fitting subsphere of one dimension lower is fit, irrespective of whether that subsphere is a great or small subsphere.

The approach begins with the data on the d -dimensional sphere. The best subsphere of $d - 1$ dimensions is fit by an analytical computation. As illustrated in




Fig. 7 Computation of subsphere, projection, and scores in Principal Nested Spheres

Fig. 7, the result is formed by (and recorded as) the pole (axis) position \mathbf{w}^1 of the subsphere on the d -sphere and its latitude (angle from the pole) ψ_1 . Then it projects the data onto the subsphere of dimension $d - 1$. The process begins anew, finding the $(d - 2)$ -dimensional subsphere best fitting the data that has been projected onto the $(d - 1)$ -dimensional sphere. The process repeats itself until a 1-sphere (a not necessarily great circle) with points on it has been arrived at. The final step is to find the mean (best-fitting 0-sphere) as the best fitting point to (geodesic mean of) the data points that have been projected onto the circle.

The mean is successively projected forward through the dimensions $2, 3, \dots, d$ using the subsphere-to-sphere transformations that have been recorded. The resulting point on the d -sphere gives the place in the data space which is the backwards mean of the shape component that was analyzed on this sphere.

As with PCA, we record the score of the j th data point relative to the projection of the subsphere of dimension $d - k + 1$, $k = 1, 2, \dots, d$, to the subsphere of dimension $d - k$ as the signed projection distance $z_{d-k+1,j}$ of that data point between those two spheres (see Fig. 7). The random variable $z_{d-k+1,\cdot}$ is a zero-mean variable that can be understood as being in Euclidean space. Thus, the $d \times N$ array Z with elements k, j being $z_{d-k+1,j}$ represents the Euclideanized, 0-mean distribution of the respective shape component of the data on the d -sphere.

5.3 Compositing component distributions into an overall probability distribution

In our method of analysis of CPNS the distribution of the 62-case s-rep data on each of the component spheres is analyzed by PNS. The compositing approach has already been described in [15]; here we detail it for analysis of s-reps. In that case, with 67 spheres, this yields 67 Z arrays. One of them (for the spoke tail points) is 61×62 and the other 66 (for the spoke directions) are 2×62 . We must now account for the correlations between the various s-rep components, including both those that live on spheres and those that live in a Euclidean space. To prepare for that, for each s-rep the Euclidean components, $k = 1, 2, \dots, 67$, each of the form log length or log scale factor, must each have their mean computed (in the ordinary way

of summing the entries across the data items and dividing by the number of entries, 62) and then subtracting that mean from each data item's corresponding component. (Strictly speaking, that data lives in a 62-dimensional space, but the SVD analysis to follow is insensitive to using the original 67 dimensions.) Call the result of this the 67×62 array $Z^{Euclidean}$.

Next each Z array row must be made commensurate; we give each entry the units of the object boundary offset produced by any row entry. This is accomplished by

1. multiplying the two rows giving z entries for each of the 66 spoke directions by the average of the length of that spoke over the training cases;
2. multiplying the n (66) rows giving entries giving a mean-subtracted log spoke length by the average of the length of that spoke over the training cases;
3. multiplying the $N - 1$ (61) rows giving z entries for the scaled and centered medial point tuple, as well as the row giving the log of the scale factor, by the geometric mean of the point tuple scale factors over the training cases.

We produce a zero-mean Euclidean representation of our data by piling all of these scaled Z arrays (67 from spherical analysis and 1 from Euclidean analysis) on top of one another, producing the 260×62 array Z_{comp} (Fig. 8), to which the same analysis as is done in conventional PCA can be applied. By this analysis, SVD (singular value decomposition) on Z_{comp} yields a list of eigenvectors (principal directions) \mathbf{v}^i and eigenvalues (principal variances) σ_i^2 . Together with the backward means for each of the spheres and the mean of the Euclidean variables, this is the PCA-like result reflecting our understanding that s-reps live on the Cartesian product of 67 spheres and a Euclidean space.

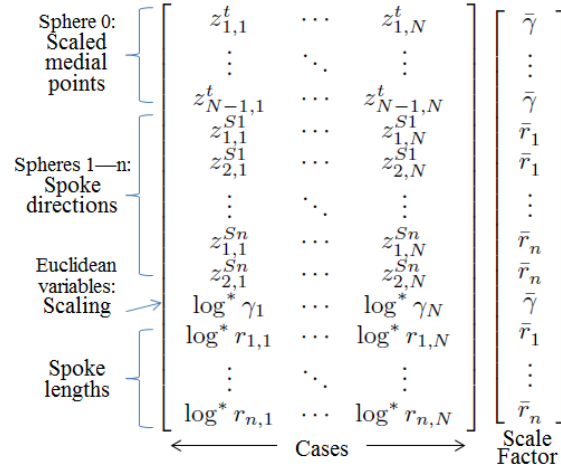


Fig. 8 Left: the composite array Z_{comp} before row scaling for N s-reps. * indicates that the variable's row mean has been subtracted. Right: the row scale factors.

The CPNS mean is formed by compositing the backward means from each of the component spheres and the Euclidean mean. Each component mean is a point on its respective sphere or Euclidean space. The mean s-rep is achieved by compositing these points into spoke directions, spoke tail points, and spoke lengths. Unlike the Euclidean mean or the forward intrinsic mean, the spherical components of the CPNS mean live not only on their spheres but on all the best fitting subspheres. It therefore can be expected that the CPNS mean is more representative of the training shapes than any of the forward means.

From the CPNS analysis, we can choose a number of eigenmodes that captures a desired fraction of the total variance. This forms a shape space for s-reps as implied by the training cases. Every s-rep in the shape space is given by a tuple $\underline{\alpha}$ of scores of the eigenmodes, with associated Mahalanobis distance $\Sigma_i(\alpha_i^2/\sigma_i^2)$. Given $\underline{\alpha}$, one can form the vector $\Sigma_i\alpha_i\mathbf{v}^i$ and then apply the row scale factors computed in preparing for the SVD. The entries in the resulting column vector can be divided into the tuples corresponding to each s-rep component, just as the columns in the array Z_{comp} were divided. Each component will have its value unscaled by the row scale factors previously discussed. The Euclidean components are then ready for adding their mean. Each group of components from a spherical space is ready for projection into its successive spherical spaces represented by the latitude angles ψ_i and polar directions \mathbf{w}^i , beginning with the backwards mean for that sphere. Compositing the results on each of the 68 spaces into spoke directions, spoke tails, and spoke lengths yields an s-rep that can be displayed or used to provide s-rep-relative coordinates for the points in the image space in which an s-rep was originally fit.

The results of a CPNS analysis on hippocampi, some of which were illustrated in Fig. 4, are described in the following section.

6 Analyses of populations of training objects

We analyze 62 left hippocampi segmented from MR images. These are from control (normal) individuals in a study of schizophrenia [26]. Each had an initial model that had been fit years ago using an earlier version of the first three stages of the s-rep fitting program; these fits were neither tight with the object boundary nor did they satisfy the spoke-to-boundary orthogonality and crest fitting objectives. These initial models were each translationally, rotationally, and scale aligned to their distance functions, and the resulting models were fit using the geometric (atom-by-atom and spoke) stages described as the second and third stages in section 3.1.2. Our objective was to compare fitting using CPNS followed by statistical analysis via CPNS with fitting using PGA followed by statistical analysis via PGA.

Fifty of the third-stage fits were chosen as satisfactory, and this set was dealt with in both ways. Fifteen eigenmodes were chosen from each set as the dominant eigenmodes. The result was two shape spaces, one with a mean and eigenmodes based on PGA on the 50 fits and one with a mean and eigenmodes based on CPNS of the same 50 fits. Following this analysis, two applications, one PGA-based and one

CPNS-based, of the fourth (shape space) and fifth (spoke) stages (see section 3.1.2) were applied to produce the two sets of fitted s-reps that we wished to compare. In each application, each of the 62 hippocampi was initialized by the mean appropriate for that set, and it was first translated, rotated, and scaled into the hippocampus distance image to be fit. That result was taken as the initialization of an optimization over the corresponding shape space; the spoke optimization of the fifth stage was applied to the result. These fits seemed generally satisfactory for all 62 hippocampi. Let us refer to the results as the PGA-based fits and the CPNS-based fits.

CPNS analysis was applied to the CPNS-based fits, and PGA was applied to the PGA-based fits. In the following we compare the eigenvalue plots, the means, and the eigenmodes between these two analyses.

In both analyses the first mode of variation dominated, and it was largely uniform scaling. It formed 70% of the total variance in the CPNS analysis and 90% of the total variance in the PGA analysis. This is consistent with a longstanding approach in neuroimage analysis of focusing on hippocampus volume as a discriminative parameter. Removing this source of variation from both the eigenvalue list and the total variance in the corresponding analysis yielded the eigenvalue and cumulative eigenvalue plots shown in Fig. 9. There is an obvious compression to the left of the eigenvalues in the CPNS plot relative to those in the PGA plot. This is quantified by looking at how many eigenmodes were necessary to achieve 90% of the remaining variance. The CPNS-based fitting and analysis requires only 7 additional modes to achieve 90% (thus 97% of the total variance when the first mode is included), whereas the PGA-based fitting and analysis requires 14 additional modes to achieve 90%. We also notice that, whereas the fitting after the first three stages, i.e., the geometry-based stages, produced analyses by either PGA or CPNS that required 20 modes each to reach 90% of the variance in modes 1 on, the amount of fitting-based noise produced by the stage optimizing over the statistical shape space followed by spoke stage refinement is so lessened that only 8 modes (the first, scaling mode and 7 more) are now necessary to reach 97% of the variance in modes 1 on.

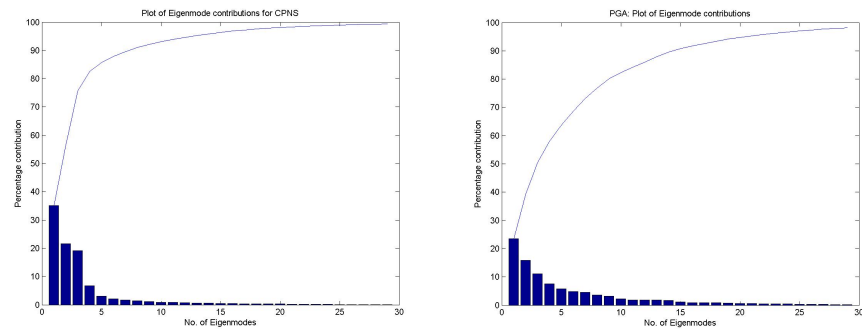


Fig. 9 Eigenvalues (bars) and cumulative eigenvalue (curve) plots as percentage of total variance for CPNS (left, requiring 7 modes for 90% of variance) and for PGA (right, requiring 14 modes for 90% of variance) after removal of the first mode.

Now we compare the CPNS-based backward mean s-rep to the PGA-based forward mean s-rep. Both means look like a hippocampus, but as seen in Fig. 10, the CPNS mean is longer, related to the fact that the PGA-based fits were noisier near the tail and head of the hippocampus.

Finally, we compare the two analyses in terms of the shape changes shown by their eigenmodes. Because the contrast is too complex to show here, we summarize what we have observed. Beyond the first, largely scaling, eigenmode, none of the following 4 eigenmodes from the PGA set qualitatively match any of the following 4 eigenmodes from the CPNS set, and vice versa. Roughly, the CPNS eigenmodes seem to capture more in the way of bendings and twistings, and the PGA eigenmodes seem to capture more in the way of rotations and small variations in eccentricity. We suggest that the reason is that CPNS by its design is more sensitive to changes in local directions (e.g., of spokes), so, relative to PGA, its preliminary shape space enables more efficient modeling of variations of that type. Moreover, its use in the final statistical analysis senses those variations more efficiently.




Fig. 10 Implied boundaries of CPNS (cyan) and PGA (yellow) means of the 62 left hippocampi of normal individuals

7 Extensions and discussion

The method of composite principal nested spheres applies to any data that are best analyzed as a composite of components lying on spheres and possibly including Euclidean components. The analyses of such data done so far have suggested that CPNS can provide more efficient analysis (fewer eigenmodes), stabler s-rep fitting, and possibly more information about twisting and bending changes than PGA and other forward methods. We have seen that fitting a geometric model, here an s-rep, to object segmentations is significantly improved by the use of a shape space based on CPNS-based analysis on preliminary fittings; the improvement is both in terms of quality of fit and in terms of making a second statistical analysis on the models informative.

Shape analysis via CPNS extends beyond s-reps to essentially every other representation of shape. Shape inherently involves directions and curvatures (derivatives on directions). As a result, it can be expected that CPNS will provide similar advantages when applied to other representations of shape. The following list some candidate representations suitable for CPNS analysis.

- Point distribution models (PDMs), made of a large tuple of boundary (and possibly other) points. On 10 lung PDMs for 10 phases of a patient's respiration, CPNS required only required 1-2 eigenmodes vs. 3 for PCA. As with s-reps, CPNS required fewer modes than the forward method, here PCA.
- Point and normal distribution models: i.e., boundary point PDMs with a boundary normal supplied with each point

- Quasi-tube s-reps
- Branching objects
- Multi-object complexes
- Warps, i.e., displacement vector images or time-indexed velocity vector images representing diffeomorphisms or other invertible deformations.
- Fiber conglomerates, such as those studied by [23].
- Tensor images
- Covariance matrices

The method of s-rep fitting followed by correspondence improvement and CPNS analysis requires a variety of improvements, listed in the following, together with testing on many more shape populations. Once verified, it seems suitable for applications in segmentation by posterior optimization [27, Chapter 9], hypothesis testing [29], and classification [25].

Improvements in s-rep fitting needed include spoke interpolations at the crest using the S_{rad} and S_E matrices [11], fittings into multi-crest distance images, and generalization of entropy-based correspondence optimization [3] to spokes.

Acknowledgements

We acknowledge the contributions of Jörn Schulz in studying the effects of global rotation, folding, and twisting on skeletal spoke vectors and for fitting s-reps to the hippocampi. We thank Anna Snyder for help in producing the reference list.

References

1. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A log-Euclidean framework for statistics on diffeomorphisms. In: Medical Image Computing and Computer-Assisted Intervention, vol. 9, pp. 924–931 (2006)
2. Blum, H.: A transformation for extracting new descriptors of shape. In: W. Wathen-Dunn (ed.) Models for the Perception of Speech and Visual Form. Cambridge, MA: MIT Press (1967)
3. Cates, J., Fletcher, P.T., M. Styner M. E. Shenton, R.T.W.: Shape modeling and analysis with entropy-based particle systems. *Inf. Process. Med. Imaging* **20**, 333–345 (2007)
4. Cootes, T.F., Taylor, C., Cooper, D., Graham, J.: Training models of shape from sets of examples. In: D. Hogg, R. Boyle (eds.) Proc. British Machine Vision Conference, pp. 9–18. Berlin. Springer (1992)
5. Cootes, T.F., Twining, C.J., Petrović, V.S., Babalola, K.O., Taylor, C.J.: Computing accurate correspondences across groups of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **32**, 1994–2005 (2010)
6. Damon, J.: Smoothness and geometry of boundaries associated to skeletal structures I: sufficient conditions for smoothness. *Annales Inst. Fourier* **53**, 1001–1045 (2003)
7. Damon, J.: Swept regions and surfaces: modeling and volumetric properties. Conf. Computational Alg. Geom. 2006, in honor of Andre Galligo, Special Issue of Theor. Comp. Science **392**, 66–91 (2008)
8. Dryden, I.L., Mardia, K.V.: Statistical Shape Analysis. Wiley, Chichester (1998)

9. Fletcher, P.T., Lu, C., Pizer, S.M., Joshi, S.: Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Trans. on Medical Imaging* **23**, 995–1005 (2004)
10. Fletcher, P.T., Venkatasubramanian, S., Joshi, S.: The geometric median on Riemannian manifolds with application to robust atlas estimation. *NeuroImage* **45**(1), S143–S152 (2009)
11. Han, Q., Pizer, S.M., Damon, J.N.: Interpolation in discrete single figure medial objects. In: *Computer Vision and Pattern Recognition (CVPR) - Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA)* (2006)
12. Huckemann, S., Hotz, T., Munk, A.: Intrinsic shape analysis: Geodesic PCA for Riemannian manifolds modulo isometric lie group actions. *Statistica Sinica* **20**(1), 1–58 (2010)
13. Huckemann, S., Ziezold, H.: Principal component analysis for Riemannian manifolds, with an application to triangular shape spaces. *Adv. in Appl. Probab.* **38**(2), 299–319 (2006)
14. Jung, S., Dryden, I.L., Marron, J.S.: Analysis of principal nested spheres. To appear in *Biometrika*. Available at <http://midag.cs.unc.edu/bibliography.html> (2012)
15. Jung, S., Liu, X., Marron, J.S., Pizer, S.M.: Generalized PCA via the backward stepwise approach in image analysis. In: J. Angeles (ed.) *Brain, Body and Machine: Proceedings of an International Symposium on the 25th Anniversary of McGill University Centre for Intelligent Machines, Advances in Intelligent and Soft Computing*, vol. 83, pp. 111–123 (2010)
16. Kendall, D.G., Barden, D., Carne, T.K., Le, H.: *Shape and Shape Theory*. Wiley, Chichester (1999)
17. Kurtek, S., Ding, Z., Klassen, E., Srivastava, A.: Parameterization-invariant shape statistics and probabilistic classification of anatomical surfaces. *Inf. Process. Med. Imaging* **22**, 147–158 (2011)
18. Leventon, M., Faugeras, O., Grimson, E., Kikinis, R., Wells, W.: Knowledge-based segmentation of medical images. In: S. Osher, N. Paragios (eds.) *Geometric Level Set Methods in Imaging, Vision, and Graphics*. Springer (2003)
19. Merck, D., Tracton, G., Saboo, R., Levy, J., Chaney, E., Pizer, S.M., Joshi, S.: Training models of anatomic shape variability. *Medical Physics* **35**, 3584–3596 (2008)
20. Oguz, I., Cates, J., Fletcher, T., Whitaker, R., Cool, D., Aylward, S., Styner, M.: Entropy-based particle systems and local features for cortical correspondence optimization. In: *Proc. IEEE Int. Sym. on Biomedical Imaging (ISBI)*, pp. 1637–1641 (2008)
21. Pennec, X.: Statistical computing on manifolds: from Riemannian geometry to computational anatomy. *Emerging Trends in Visual Computing* **5416**, 347–386 (2008)
22. Saboo, R., Niethammer, M., North, M., Pizer, S.M.: Anti-aliasing discretely sampled object boundaries using fourth-order Laplacian of curvature flow. In preparation for submission, available at <http://midag.cs.unc.edu/bibliography.html> (2011)
23. Savadjiev, P., Campbell, J.S.W., Descoteaux, M., Deriche, R., Pike, G., Siddiqi, K.: Labeling of ambiguous sub-voxel fibre bundle configurations in high angular resolution diffusion MRI. *NeuroImage* **41**(1), 58–68 (2008)
24. Schulz, J., Jung, S., Huckemann, S.: A collection of internal reports submitted to UNC-Göttingen study group on s-rep change under rotational transformations (2011)
25. Sen, S.K., Foskey, M., Marron, J.S., Styner, M.A.: Support vector machine for data on manifolds: An application to image analysis. In: *Proc. IEEE Int. Sym. on Biomedical Imaging (ISBI)*, pp. 1195–1198 (2008)
26. Shi, X., Styner, M., Lieberman, J., Ibrahim, J.G., Lin, W., Zhu, H.: Intrinsic regression models for manifold-valued data. In: *Medical Image Computing and Computer-Assisted Intervention*, vol. 5762, pp. 192–199 (2009)
27. Siddiqi, K., Pizer, S.: *Medial Representations: Mathematics, Algorithms and Applications*. Springer (2008)
28. Sorensen, P., Lo, J., Petersen, A., Dirksen, A., de Bruijne, M.: Dissimilarity-based classification of anatomical tree structures. *Inf. Process. Med. Imaging* **22**, 475–485 (2011)
29. Terriberry, T., Joshi, S., Gerig, G.: Hypothesis testing with nonlinear shape models. *Inf. Process. Med. Imaging* **19**, 15–26 (2005)
30. Wang, H., Marron, J.S.: Object oriented data analysis: sets of trees. *Ann. Statist.* **35**(5), 1849–1873 (2007)
31. Yang, J., Staib, L., Duncan, J.: Neighbor-constrained segmentation with 3D deformable models. *Inf. Process. Med. Imaging* **18**, 198–209 (2003)