

Object Shape Representation via Skeletal Models (s-reps) and Statistical Analysis

Stephen M. Pizer¹, Junpyo Hong¹, Jared Vicory¹, Zhiyuan Liu¹, J.S. Marron¹,
Hyo-young Choi², James Damon², Sungkyu Jung², Beatriz Paniagua², Jörn Schulz²,
Ankur Sharma², Liyun Tu², Jiyao Wang²

Skeletal models that are structurally medial provide effective object representations. This is because they include not only locations but also boundary directions and object widths. We present a skeletal object representation that we call “quasi-medial” because geometric properties associated with Blum’s [1967] medial axis are relaxed to allow the skeleton to have a pre-specified amount of branching and thus to support statistical analysis. We call this form of object representation the *s-rep*. We explain how such models can be automatically determined from object boundary data in a way that a) avoids boundary noise, b) implies a boundary that closely fits the input boundary, and c) well recognizes shape correspondences across cases. We also explain how to use Riemannian geometry to estimate probability distributions from a sample of s-reps and to find ways to classify an s-rep between two categories as trained from s-reps in each class. Finally, we describe various evidence that shows the relative strengths of s-reps vs. other object representations; we also discuss shortcomings of s-reps.

1. Introduction to Skeletal Models

There are many ways to represent an anatomic object or object tuple, entities in 3 dimensions. These representations can be categorized into those that describe the object in terms of its geometric features and those that describe it in terms of a deformation of space from an atlas containing the object. Here we will be describing a representation via geometric object features.

Which geometric features are useful to capture in the representation? Many describe only boundary geometry, given by the boundary locations and/or boundary normal directions and/or boundary curvature features. Boundary normal directions and curvatures have been shown to be particularly important shape properties. But for the many objects that have the form of a bent slab there is one additional set of features that intuitively seems important: *object widths* as they vary along the slab. The skeletal models are the only ones that explicitly capture object widths as well as boundary locations and normal directions. As laid out later in this chapter, comparisons of models according to two capabilities suggest that skeletal models are particularly strong; these capabilities are their power for classification and in providing prior probabilities to support segmentation.

¹ Principal co-authors of this chapter

² Contributors of important material to this chapter

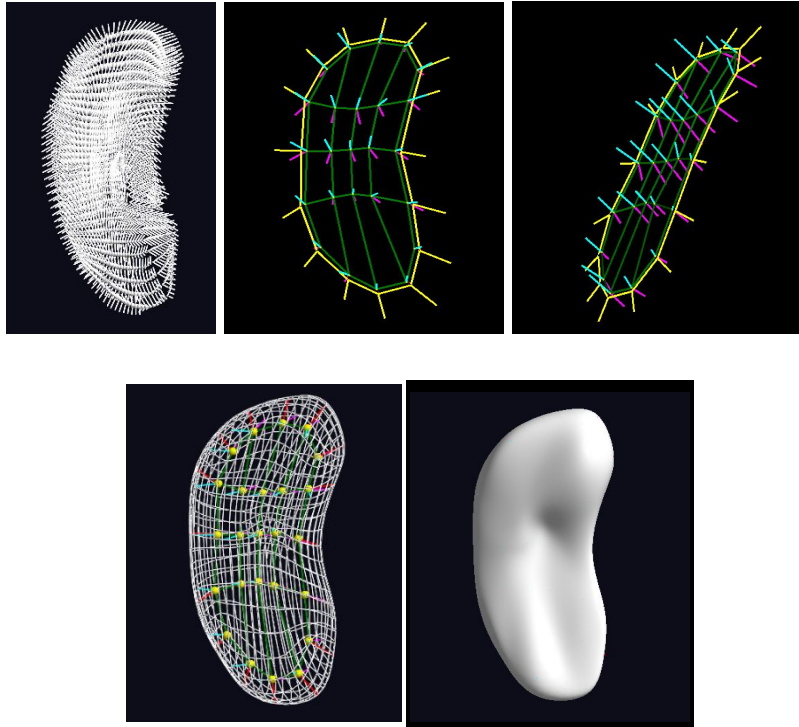


Figure 1. A skeletal model for a hippocampus. a) Densely sampled spokes. b, c) The computer representation of an s-rep, coarsely sampled spokes, from two points of view. The grid of green lines connect vertices whose locations are interior skeletal points, each repeated to be on both sides of the folded skeletal surface. The yellow lines connect vertices whose locations are on the skeletal fold. The line segments proceeding from the skeletal fold points are fold spokes. The magenta line segments proceeding from the interior skeletal points on one side of the object are interior spokes on that side. The cyan line segments proceeding from the interior skeletal points on the other side of the object are also interior spokes. d) The computer representation and the s-rep's implied boundary shown as a wire mesh. e): The boundary implied by the s-rep.

As illustrated in Figure 1, skeletal models for objects are best understood as formed by a skeleton with the same topology as the object boundary, together with a vector function on the skeleton with the following properties: the tail of each vector is on the skeleton and the tip of each vector is on the object boundary, and the set of all such vectors have no crossings within the object and fill the object interior. We call these vectors “spokes”. The skeleton is a collapsed form of the object; it is folded onto itself such that except at the folds, at least and typically two places on the skeleton are at the same location. That is, the two spokes emanating from these places share tails; together each spoke pair captures the object width there. Also, the spoke tips cover the whole boundary and the spoke directions capture boundary directions. While for both 2D and 3D objects such skeletons exist and are useful, we focus here on the ones in 3D.

The earliest form of skeletal models was “medial”. While there were many variants (see Chapter 1 of [Siddiqi & Pizer 2008]), the one due to Blum [1967] was largely settled upon, and the mathematics of the Blum medial loci were deeply developed (esp. see

Chapters 2 and 3 of [Siddiqi & Pizer 2008]). The basic properties of the Blum representation are that

- 1) For slab-like objects most skeletal points have two spokes whose tails are at that point and for which the two spokes have equal lengths, making the skeleton exactly centered between two opposing boundary points and making the common length a half-width of the object;
- 2) At the tip of each spoke it is orthogonal to the object boundary – it captures the boundary position and normal.
- 3) Skeletal branches occur where three spokes share a tail point. The topology of the skeleton of 3D objects also has more complicated skeletal points (see chapter 2 by Giblin & Kimia of [Siddiqi and Pizer 2008]).

However, when the input boundary has detail, the Blum skeleton (“medial axis”) is extremely bushy (Figure 2). The variability in this branching across cases makes that object representation very challenging for object shape statistics. Moreover, pruning the representation was found to be fraught [Siddiqi and Pizer 2008], and the locus of spoke tips from the pruned axis containing only intuitively essential branching could not succeed in giving a reasonably accurate approximation to an input object’s boundary.




Figure 2. Left: A bushy Blum medial axis of a rectangle with boundary noise (from G. Székely, chapter 6 of [Siddiqi and Pizer 2008]). The desired medial axis consists only of the five blue line segments shown in the right figure within the noise-free rectangle; the remainder of the lines shown interior to the noisy rectangle are the unwanted bushy components of the computed axis.

Thus Damon and Pizer [Pizer 2013] developed a more general, more flexible form of skeletal model and its mathematics [Damon Chapter 3 in Siddiqi and Pizer 2008]. This model was named the “s-rep”; this chapter describes the s-rep as well as its mathematics and statistics. By only penalizing, rather than preventing, deviations from the medial properties of a) precisely matching the input boundary, b) having spokes that are precisely orthogonal to the input boundary, and c) having interior spokes that have exactly the same length as their partners proceeding to the other side from the same skeletal point, the s-rep could imply a boundary that accurately fits the input boundaries while containing only essential branching. Indeed, so many anatomic objects could be represented with no branching whatsoever that only limited work has been done to describe essentially branched objects with s-reps.

For a population of instances of an object, it is desirable to have correspondence between interior locations across instances. This requires first, the same branching structure, second, correspondence of the fold curves, and third, correspondence within the respective skeletal sheets.

While the discussion so far has been for slab-like objects that have the topology of a sphere, objects with other topologies can be also usefully represented with skeletal models. In particular, objects that have the topology of a cylinder without its top and bottom, sometimes called “generalized cylinders”, are also usefully represented skeletally (Figure 3). In these objects the skeleton is formed by collapsing the quasi-tubular structure to a space curve that is infinitesimally dilated to a tube. In this structure, here again, there is one spoke per skeletal point and the spokes fill the object and do not cross in the object interior.

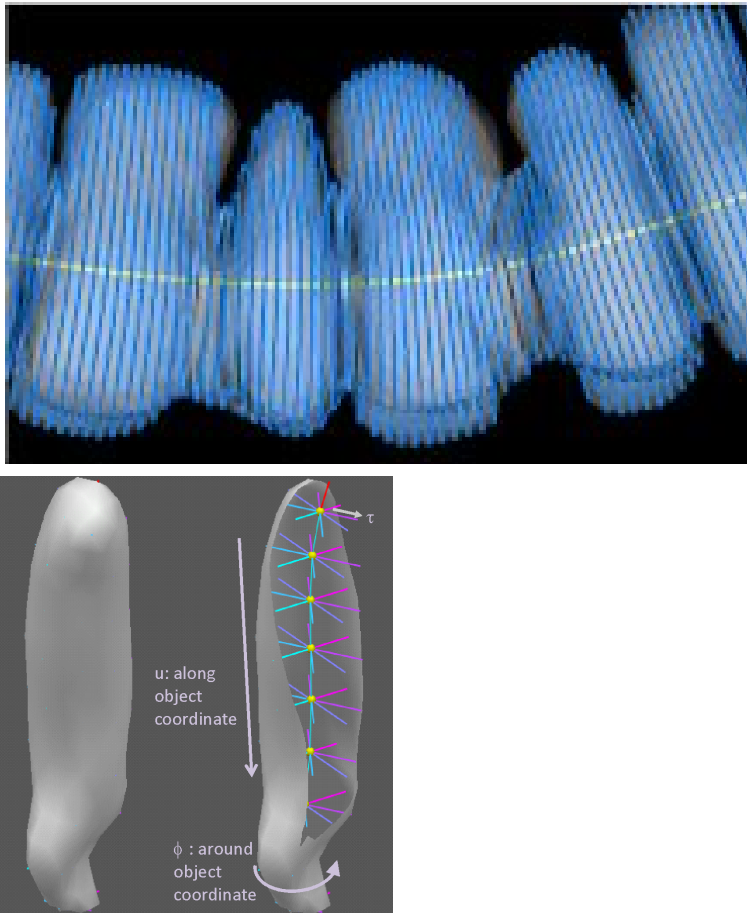


Figure 3: Left: A section of the large intestine represented as a generalized cylinder [Ma 2018]. The center curve is the axis of the generalized cylinder, and the blue curves orthogonal to the axis show the cross-sections of the generalized cylinder. Right: A discrete s-rep with tubular topology, implying the boundary shown in the left half of the panel.

2. Computing an s-rep from an image or object boundary

The almost innumerable algorithms for transforming an object boundary into its Blum medial axis fall into a few categories, summarized in Chapters 4-7 of [Siddiqi and Pizer 2008]. There you will find an algorithm in each category that the authors deemed to be the best at the time that book was written. The categories are 1) computing shocks of a grassfire boundary evolution; 2) computing skeletons that are lists of voxels using discrete methods, e.g., via the distance transform; and 3) using the Voronoi diagram to compute skeletons. These algorithms face in common the problem that the function from object boundaries to skeletal loci is ill-conditioned; that is, it turns small errors into large ones and in particular produces bushy skeletons that need pruning. This is problematic because of the failure to have found adequate methods for pruning.

What is needed are algorithms for finding a skeleton that recognize only the specific branching of the object that is essential to its shape. For example (Figure 4), a tube with a smooth bend going around 180 degrees to become parallel close to itself would have a single curvilinear skeleton unlike the medial skeleton that has a branch at the bend.




Figure 4: Bent bar with branching medial axis. Left: with intuitive axis; right: with approximate Blum medial axis.

One algorithm for extracting a quasi-medial skeleton with this trait works on images and generates non-branching skeletons by using the properties that height ridges generically do not branch [Damon 1999]. These images can be those with high contrast at object boundaries or can be signed pseudo-distance images, in which the voxel values are zero at the boundary, negative on the inside of the object, positive outside it, and monotonic in magnitude of distance from the object boundary. The idea is to consider probing the image via a primitive (Figure 5) that responds strongly when it is consistent with medial behavior. Specifically, the probe is a function of image position (the candidate medial position \underline{x}), spoke length (r), and the directions of two spokes ($\underline{\theta}_1, \underline{\theta}_2$); thus in 3D images the probe is 8-dimensional. The probe measures the evidence that at the two spoke tips the image is boundary-like in the spoke directions. Applying the probe produces a scalar measurement $M(\underline{x}, r, \underline{\theta}_1, \underline{\theta}_2)$ in an 8-dimensional non-Euclidean Riemannian manifold. In this algorithm a medial locus of dimension k (e.g., 2 for a slab-like object or 1 for a generalized cylinder) is a k -dimensional height ridge of M called a “core”.

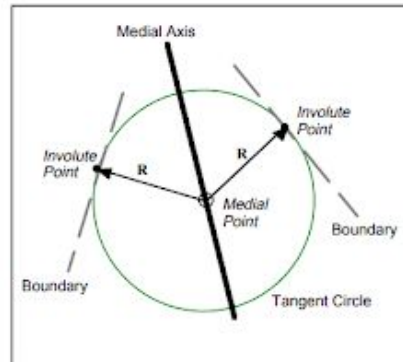


Figure 5: Medial strength probe (copied from [Katz 2002]). \underline{x} is the position labeled “Medial Point”. $\underline{\theta}_1$ and $\underline{\theta}_2$ respectively give the directions of the two spoke vectors of length R (called “ r ” in the text) proceeding from the medial point to the involute points, at which the object boundary is sensed with tangent direction (the dashed lines) perpendicular to the vectors.

These cores [Morse 1994] proved to be particularly powerful for finding blood vessels in noisy 3D medical images [Aylward 2002; Fridman 2004]. Fridman also showed they could support finding blood vessel branches by exploring the behavior of the function M , and he demonstrated this property as well for slab-like branches on slab-like objects.

However, cores had several problems. The algorithms for computing them worked by ridge following, and like all such algorithms, initializing the ridge was not easy and the ridges could break in areas where the medial evidence is weak, where they would change into height saddles.

Reflecting that transforming boundaries into skeletal structures mapped small errors into large ones, it became clear that the inverse process (skeletal model mapped to the boundary) would be far preferable, as it would map large errors into small ones. Thus the approach that we have adopted was born: *fitting* a skeletal model with fixed branching topology into boundary data. This has proven to have adequate stability to allow a variety of uses, presented later in this chapter.

The s-reps that we invented have developed through a number of computer representations. In the medial literature it has been traditional to think of the skeletal surface of an object as a bounded surface (in 3D; a curve in 2D) with two spokes from all locations except for branch and surface-end locations; in early incarnations of the s-rep, this was how it was represented in the computer. However, as illustrated in Figure 6, it is notably clearer and more generalizable to think of the skeleton as being a folded surface produced by collapsing the two sides of an object onto each other, with each location on that folded surface having a *single* spoke. In this view the end points in the traditional view have become the fold. Thus, other than at the fold, two skeletal points are at the same location, but the skeleton has the same topology as the object being represented. Mathematically, the s-rep is written as follows: Let \underline{u} parametrize the boundary of the

object being represented. Then the s-rep $m(\mathbf{u}) = (\underline{p}(\mathbf{u}), \mathcal{S}(\mathbf{u}))$ is also parametrized by \mathbf{u} , where $\underline{p}(\mathbf{u}) = (x(\mathbf{u}), y(\mathbf{u}), z(\mathbf{u}))$ is the skeleton and $\mathcal{S}(\mathbf{u})$ is the spoke emanating from the skeletal point $\underline{p}(\mathbf{u})$. Moreover, $\mathcal{S}(\mathbf{u}) = (\mathbf{U}(\mathbf{u}), r(\mathbf{u}))$, where the unit vector \mathbf{U} is the spoke direction and r is the spoke length.




Figure 6. An ellipsoid and its skeleton, formed by collapsing the ellipsoid with respect to its principal axis of shortest length. The doubled ellipse shown by the top side in orange and the bottom side in yellow actually share locations; the slight offset is used in this visualization to allow you see the two ellipses.

In the computer representation the skeleton and its spokes are spatially sampled (Figure 1 – in it strictly the dense spokes are interpolated from the coarse spokes in the computer representation) into a sparse version that we also refer to as “the discrete s-rep”. Such a computer representation implies a continuous s-rep by the use of a spoke interpolation method supported by the mathematics of skeletal structures due to Damon (chapter 3 in [Siddiqi and Pizer 2008]) – see Section 3.

Paul Yushkevich [2006] developed a useful alternative form in which the interpolation is built in through the use of an explicit analytic representation; this is done via the use of a specialized spline on (x, y, z, r) . Like all splines, this one is represented by spline coefficients that can be understood as control “points”. While this form of skeletal representation of 3D objects has found important medical application where physical modeling of anatomic objects is desirable, for statistical applications it has not been shown to be competitive with our s-reps based on skeletal sampling.

3. Skeletal Interpolation

To fit an s-rep to an object boundary, the spokes must be dense enough. Thus, the skeletal surface and the spokes from it must be interpolated from the sparse representation that we use (Figure 1). Skeletal interpolation is also needed to display the boundary implied by the s-rep. Spokes on a skeletal surface follow the skeletal geometry of each spoke direction living abstractly on a 2-sphere and each log of spoke length living in a Euclidean space. That is, spokes live in a curved abstract space. Moreover, the fact that the s-rep is approximately medial places soft constraints on that space. We have shown that actually using that geometry is necessary to ensure that the interpolated s-rep is usable.

It would probably be mathematically best to interpolate the skeletal surface and the spokes in a single interpolation. However, for simplicity Vicory [Tu, Vicory 2016]) divided this into a skeletal surface interpolation operation and a spoke interpolation on the interpolated skeletal surface.

The skeletal surface interpolation method uses standard polynomial-based methods. One idea would use the property that the normal to the skeletal surface can be expected to be in a direction near the difference of the spoke directions; this is a fact in the medial case. Since our skeletal representations are only quasi-medial, this would have to be done via a penalty, i.e., soft constraint, on the angles of the paired spokes to the skeletal surface being computed, and to date we have not used this approach.

Our method for spoke interpolation depends on first interpolating the spoke directions U and then interpolating the spoke lengths r . Our experience shows that this tack is preferable because we have found a good way to interpolate the directions and have found that the overall result is then very robust against approximation errors in the spoke length interpolation.

Let the skeletal surface be parametrized by (u_1, u_2) , where both parameters are integers at the corners of a quadrilateral in the grid on which the discrete s-rep is specified. Thus, the discrete s-rep gives both r and U at these quadrilateral corners. Consider interpolation of the spoke directions U at any point $p(u_1, u_2)$ within any grid quadrilateral on the skeletal surface. Our plan for interpolation of r is based on a 2nd-order Taylor series, for which we need not only the spoke directions U but also their first and second order derivative values U_{u_i} and $U_{u_i u_i}$ for $i=1, 2$ at arbitrary points in the quadrilateral. Spoke directions live on the unit 2-sphere S^2 . Thus, the sort of finite difference calculations that must be used in order to compute U at our discrete skeletal points should be done on the sphere. These calculations are done by representing the discrete spokes U as unit quaternions and thus its derivatives with respect to u_i as derivatives on the sphere. Using these derivatives, Vicory applies the *squad* method of interpolating quaternions to estimate the spoke direction U at an arbitrary point interior to a quadrangle of discrete points by fitting Bezier curves to the quaternions on the surface of the sphere. This approximation allows the computation of not only the U values but also their directional derivatives of both first and second order in either u_1 or u_2 .

Given the ability to evaluate U and its derivatives in a quadrilateral, we need to interpolate the r values in a way consistent with skeletal geometry. Spokes can be written $S=rU$. The derivatives of the spoke at a skeletal location p with respect to a step in direction v in either of the two orthogonal directions u_1 or u_2 must follow $S_v = rU_v + r_v U$, from which it follows that $r_v = S_v \cdot U$. Also, $S_{vv} \cdot U = (S_v)_{v \cdot} \cdot U = r_{vv} + r U_{vv} \cdot U$. From this a Taylor series in the length d of a small step in direction v together with three forward-distance derivative approximations yields $r(p + dv) = \frac{1}{2} (S(p) + S(p + 2dv)) \cdot U(p + dv) - \frac{d^2}{2} (S(p) \cdot U_{vv}(p))$. Because the same

³ Subscripted variables refer to directional derivatives in the direction of the subscript

mathematics works using a Taylor series in the backwards direction about $p + 2dv$, for symmetry and to reduce approximation error the results of the two versions should be averaged, yielding the final formula as

$$r(p + dv) = U(p + dv) \cdot \left(\frac{1}{2} (S(p) + S(p + 2dv)) \right) - \frac{d^2}{4} (S(p) \cdot U_{vv}(p) + S(p + 2dv) \cdot U_{vv}(p + 2dv)) .$$

This formula allows computing the spoke half-way between two horizontally adjacent quadrilateral corners and using successive subdivision with it by halving as many times as necessary to get the desired small spacing in u_1 . Applying the same method separably (on the results) in the u_2 direction, yields a spoke at any successively subdivided point within the quadrilateral. Finally, since the method gives different results when you apply it first in u_2 and then in u_1 , we compute using both orders and average the results.

At a skeletal fold, the skeletal surface's lack of smoothness prevents direct application of the aforementioned method. We solve this problem by first dilating the fold curve into a tube with very narrow radius. We then find where the spoke at the fold curve intersects the tube. Then we use the method for smooth surfaces to compute the continuation of the spoke from the smooth surface of the tube to the object boundary.

4. Skeletal fitting

After trying many approaches, the best ways we have found to fit an s-rep to the given boundary of a particular object is to initialize a process with a reference s-rep appropriate to the object, to map that s-rep via boundary correspondences to the target, and then to complete the process by a refinement of the mapped s-rep.

Two methods for initialization have been created. The newer form [Hong 2018] involves using an ellipsoidal reference object; this is attractive both because no user-chosen reference object is needed and because the ellipsoid has a mathematically known unbranching *medial* representation, with a folded flat ellipse as the skeletal surface. The method uses mean-curvature flow on the input object boundary in small steps until it approaches closely enough to an ellipsoid, from which a sampled s-rep can be derived. The small steps allow the establishment of correspondence between points on the boundary before a step and those after a step. Then the boundary correspondences for each step, together with correspondences at the skeletal fold can be used in reverse to obtain a thin-plate-spline transformation from the after-flow state to the before-flow state. The skeletal fold correspondences are computed from the normal and the radius of curvature at the corresponding crest points. The small-step transformations can be composed and applied to the ellipsoidal s-rep's spoke tips (boundary points) and tails (skeletal points) to yield an initialized s-rep in the original object.

This method is new, so it has been shown to work for only a few objects, namely the relatively blob-shaped hippocampus (see Figure 7), the thin-width lateral ventricle, and the narrowing shaped caudate nucleus (for which a special primitive to handle the cusp-end has been designed [Hong 2018]), the parotid gland, and the rather complicated-shaped mandible. Objects with sharp points or crests, such as the mandible, need special attention to avoid some of the mesh's tiles becoming too small. Alternative means of mapping the ellipsoid's s-rep back to the original object should perform faster.

This method establishes a correspondence among a population of objects as follows. After the mean curvature flow each of the objects is mapped to an ellipsoid, the geometric means of the corresponding principal radii can be computed, and each object's ellipsoid can be mapped to the mean ellipsoid by the implied expansion of its three principal radii. Thereby, the sampled s-rep of the mean ellipsoid (Figure 7) can be mapped to each object, after which each object's sampled s-rep can be deformed back to its original object.




Figure 7. The skeletal model derived from a hippocampus boundary via the mean curvature flow method. a) The ellipsoid s-rep from which the hippocampus s-rep is derived, together with a wire-mesh showing the ellipsoid, and b) its implied boundary. c,d): The fitted hippocampus s-rep from two points of view. The interior skeletal points are grid vertices in green. Magenta and cyan spokes emanate from them. The yellow curve is the fold, with fold spokes emanating from it. e) The implied boundary of the s-rep after initial fitting. f) The implied boundary of the refined s-rep. g) The refined s-rep showing both the implied boundary as a wire mesh and the s-rep itself.

The older form of initialization involves producing s-rep fits for many instances of an object and then taking the reference model to be their backwards mean, as described in Section 6. In this form we have found that a deformation of the reference object boundary into the target object boundary can be calculated by the thin-shell demons method [Zhao 2015] (preceded by a similarity transformation), which yields a transformation of the boundary locations with correspondence that can be generalized to 3D using thin-plate splines and then applied to the tails and tips of the reference s-rep.

For both methods of initializing an object's s-rep, refinement of the initialized model is done by optimization of a penalty measuring the deviation of the s-rep from the medial ideal. The penalty is the integral over the spokes of a weighted sum of three terms, one measuring the fit of each original and interpolated spoke tip to the target object boundary, another measuring how orthogonal the spoke tip is to the boundary, and the third measuring how mathematically proper the model is at that spoke tip, i.e., how stable the model is against spoke crossing in the object interior. The measure of fit integrates $|d(\underline{p}(\mathbf{u}) + r(\mathbf{u}) \mathbf{U}(\mathbf{u}))|$ over the spokes, where $\underline{p}(\mathbf{u}) + r(\mathbf{u}) \mathbf{U}(\mathbf{u})$ is the location of the spoke tip and d is the distance of that location from the object boundary. The measure of spoke-to-boundary orthogonality is 1 minus the cosine of the angle between the spoke

direction and the boundary normal, i.e.,
 $1 - [U(\mathbf{u}) \cdot \nabla d(\underline{p}(\mathbf{u}) + r(\mathbf{u})U(\mathbf{u})) / |\nabla d(\underline{p}(\mathbf{u}) + r(\mathbf{u})U(\mathbf{u}))|]$.

The measure of stability against spoke crossing heavily penalizes spoke crossing and slightly penalizes near-crossing of spokes. In particular, the spoke-crossing test devised by Damon (Chapter 3 in [Siddiqi and Pizer 2008]) for each spoke in the continuous skeletal model, namely $1 - \det(rS_{rad}) > 0$, is used as the argument of the penalty weight, where S_{rad} is what Damon calls the *radial shape operator*. This operator is analogous to the familiar shape operator from differential geometry, which is a 2×2 matrix that when applied to a small step $\Delta \mathbf{B} = a\mathbf{v} + b\mathbf{w}$ on a surface B whose tangent plane is spanned by the frame (\mathbf{v}, \mathbf{w}) yields the change of the surface normal for that step, also represented in the frame (\mathbf{v}, \mathbf{w}) . Analogously, the radial shape operator is also a 2×2 matrix that when applied to a small step $\Delta \underline{p} = a\mathbf{v} + b\mathbf{w}$ on the skeletal surface yields the component of the spoke direction change ΔU in the skeletal surface for that step, also represented in the frame (\mathbf{v}, \mathbf{w}) . The S_{rad} matrix is computed from derivatives in the orthogonal directions \mathbf{v} and \mathbf{w} , respectively, of the spoke direction U ; as described in Section 3, this is computed by quaternion derivatives on the sphere.

In some cases it is useful to combine the spoke crossing penalty with a penalty for non-uniformity of the discrete spoke positions on the skeletal surface. In particular, the penalty can be monotonic with the size of the difference between each interior discrete skeletal position and the average of those positions at adjacent grid positions.

5. Correspondence

In producing statistics on any geometric representation of a training population, and in particular for s-reps, correspondence of the primitives is important. For s-reps two methods for producing correspondence have been developed. The first depends on starting the fitting of each object in the population from the same reference model or related reference models, such as the ellipsoids in the method described in Section 4. This already produces useful correspondence. Somewhat improved correspondence, according to a measure of entropy among the object s-reps (see Section 7d), can be obtained by the method developed by Tu et al. [2016], in which geometric entropy is maximized (the probability distribution is tightened) with an additional entropy term in the objective function that encourages uniformity of the spokes within each object in the population.

Tu (Figure 15 below) showed that when correspondence optimization was applied measures of specificity and of generalization on s-rep-implied boundary points (see Section 7c) were improved. But also there was an improvement as compared to the method of Cates et al. [2007], where the boundary points' entropy is directly optimized by the same entropy-based method that the s-reps method was modeled after. We take this improvement as a value-added of including features of object width and spoke direction in addition to boundary point location.

6. Skeletal statistics

The major scientific challenges in modern medical image analysis arise at the population level and are addressed using random samples drawn from the population of interest. For a given skeletal configuration (discrete s-rep), the set of all possible skeletal models is called the *shape space*, and the sampling process is understood via a probability distribution on that space. Statistical analysis is then used to infer underlying population properties from the available sample.

Because, as described in Section 6.1c, the s-rep shape space is naturally represented as typically hundreds of Euclidean variables (3D locations of atoms plus log spoke lengths) and of angles on respective spheres S^2 , sufficiently general probability distributions are most naturally considered as being defined on the corresponding Riemannian manifold (or even a Finslerian manifold).

There are a number of approaches to statistical analysis of data objects lying on a Riemannian manifold, which have led to varying degrees of success. One class of these has been called *extrinsic*; see [Patrenganaru and Ellingson 2015] for a good introduction and overview. The key idea is to treat data on the manifold as points in the embedded Euclidean space, do statistical analysis there, and project the results back to the manifold. Another class of approaches is *intrinsic* analysis, which aims to more directly incorporate the curved manifold nature of the data space into the analysis. Here we discuss several intrinsic approaches that are commonly used in the analysis of s-rep data. For more detailed discussion in this direction, see [Pizer and Marron 2017].

The intrinsic statistical methods that we will discuss all involve Gaussian probability distributions. Some methods try to apply Gaussians directly on the manifold; we will call these *directly Gaussian methods*. Others map the manifold to another Euclidean manifold in which Gaussians are used; we will call these *Gaussian via Euclideanization methods*. For example, in producing statistics on LDDMM deformations understood as geodesic paths according to a complicated metric on a very high dimensional manifold, one might try to accomplish this directly or try to Euclideanize by mapping each path to a still very high but much lower dimensional space of momenta and then do Gaussian analysis on the space of momenta, which is arguably a cotangent space [Vaillant 2004].

Our methods of statistics on s-reps fall into the category of Gaussian via Euclideanization methods.⁴ They work by using PCA-like methods to produce scores that frequently can be treated as measured via a Euclidean norm, then doing Gaussian analysis on the space of these scores, and finally mapping the results of the Gaussian analysis back onto the original manifold.

⁴ The following is appropriately modified material from our previous paper [Pizer & Marron 2017], reprinted here by permission of the copyrigher.

We first discuss methods for estimation of the probability distribution on s-reps, and later we discuss methods for classification via s-reps.

6.1 Probability distribution estimation

a. Tangent plane statistical analysis methods

The first generation of methods are intrinsic analogs of PCA for the analysis of Riemannian manifold data. A manifold is a surface in the ambient space that is smooth in the sense of having an approximating hyperplane (in the sense of shrinking neighborhoods) at every point. In this spirit, [Fletcher 2004] first proposed a directly Gaussian method that he called *Principal Geodesic Analysis* (PGA). Here ordinary PCA, which we now call forwards PCA (see our later discussion of backwards PCA), is generalized to the manifold via geodesics. The 0-dimensional approximation is the point on the manifold minimizing the sum of squared geodesic distances from the sample data points to it (the Fréchet mean); the 1-dimensional approximation is the curve on the manifold through the Fréchet mean minimizing the sum of squared geodesic distances from the sample data points to it; etc.

Because computing these geodesic distances is often unwieldy or worse, [Fletcher, Lu, Pizer, & Joshi 2004] proposed an approach using Euclideanization. This is based on the plane that is tangent at the Fréchet mean (Figure 8). The data on the surface of the manifold are represented as points in the tangent plane using the *Log map*. PCA is then performed there, and the resulting eigenvectors and summarized data are mapped back into the manifold using the *Exponential map*. The corresponding scores give a type of Euclideanization.




Figure 8. Tangent plane statistical analysis (by permission from X. Pennec)

While much useful shape analysis has been done using tangent plane methods, large gains in statistical efficiency have been realized through the development of more sophisticated Gaussian via Euclideanization methods. The reason for this can be understood from the S^2 example shown in Figure 9. The blue points are a sample of S^2 directions for a single spoke in the bladder-prostate-rectum simulator model of [Jeong 2009]. Note that these data points tend to lie along a circle, so their Fréchet mean lies

near the center of the circle. Thus, the data appear as a circle when they are mapped into the tangent plane using the Log map.




Figure 9. Set of locations of a single spoke generated by an s-rep simulator model. Shows that important s-rep modes of variation follow small circles, whose variation is poorly captured by tangent plane analysis. The orange point is the Fréchet mean, which captures the center of the data less well than the red point, which is the backwards mean (see Section 6.1b).

Hence, while the data are essentially *one-dimensional* in nature (since they just follow a single small circle), the PGA requires *two* components (because the projections again follow a circle) to appropriately describe the variation in the data. A data representation living in two dimensions is much less statistically efficient than a one-dimensional representation, in the sense that modeling with higher dimensional probability distributions is less efficient. This has motivated a search for more statistically efficient approaches to the statistical analysis of data lying on a manifold.

The first of the methods recognizing that the data is on a sphere and principally lies on a subsphere of that sphere is *Principal Nested Spheres* (PNS), motivated and described in Section 6.1b. Extension of this to more complicated manifolds, such as the polyspheres central to s-rep shape representations is given in Section 6.1c. Section 6.1d discusses a yet more efficient approach to polysphere analysis involving a high-dimensional spherical approximation followed by a PNS analysis.

b. Principal Nested Spheres

In the case of data lying in a high dimensional sphere S^k embedded in \mathbb{R}^{k+1} , a useful intrinsic version of PCA is Principal Nested Spheres (PNS), proposed by [Jung et al. 2012]. The central idea is to iteratively find a nested (through dimension) series of subspheres, each of which provides an optimal fit to the data (Figure 10). In particular, at each step the dimension of the approximation is reduced by 1, finding the subsphere that best fits the data in the sense of minimum sum of squared residuals, measured using geodesic arc length along the surface of the sphere. In the case of S^2 , as shown in figure 9, PNS corresponds to the use of a polar coordinate representation carefully chosen with its pole at the correct place on the sphere.

The signed residuals are also saved as PNS scores for that component. The concatenation of these scores, over the dimensions, becomes the PNS Euclideanization of each data point. The advantages of this approach are statistical efficiency (in terms of a lower-dimensional scores representation, as illustrated in Figure 9) and tractability, since each lower dimensional manifold is determined by the imposition of a single (usually easy to find) constraint. Note that when the PNS scores are modeled as a Gaussian distribution, the induced probability distribution on the shape manifold will keep the property of unimodality and will generally have a roughly Gaussian shape as long as the spread along the manifold is not too large.




Figure 10. An optimal small subsphere, geodesic distances ($Z_{d,1}$) forming scores, and projections onto the subsphere (by permission from S. Jung)

One reason that PNS was an important statistical landmark is that it motivated the more general idea of *Backwards PCA* as a general paradigm for finding principal components in non-Euclidean data contexts. The full generality of this idea can be found in [Damon & Marron 2014]. A key concept is that the general utility of backwards PCA follows from thinking of PCA in terms of a nested series of constraints. Backwards tends to be easier to work with because from that viewpoint, the constraints can be found sequentially, instead of needing to know the full set and then sequentially relaxing them. As noted in [Damon and Marron 2014], this idea is seen to generate (or to have the potential to generate) useful analogs of PCA in a variety of other non-Euclidean settings such as on other manifolds, for Nonnegative Matrix Factorization, and for Manifold Learning.

The example shown in figure 9 also shows that the Fréchet mean, shown as the orange point in the figure, can be a poor choice of center in that it may not be representative of the data in the sense that it lies outside the cloud of data points. A more representative notion of center is the *backwards mean*, shown as the red point in the figure. This backwards mean is an intrinsic mean that is much more representative of the data in that example than the Fréchet mean. Generally this is computed by taking one more step in PNS. In this last step the best fitting point, i.e., 0-dimensional sphere, to the S^1 approximations of the data, is found. In particular, the backward mean is the Fréchet mean of the rank 1 circular representation of the data, which can then be viewed as the best backwards rank 0 approximation. So in figure 9 the backwards mean falls on the

circle determined by the data whereas the forwards mean (the Fréchet mean) falls far from that circle.

A fundamental observation of [Pennec 2015] is that both the backwards methods (e.g., PNS), which were explicitly described in [Damon & Marron 2014], and forwards methods such as PGA rely upon greedy sequential searches (in opposite directions). That motivated a simultaneous fitting approach called *barycentric subspaces*, which are affine combinations (a generalization of Fréchet, Karcher or exponential weighted averages where negative weights allow appropriate extrapolation) of a set of *reference points*. These methods are described in detail elsewhere in this book.

c. Composite Principal Nested Spheres

[Pizer, Jung, Goswami, Vicory, Zhao, Chaudhuri, Damon, Huckemann & Marron 2013] proposed extending PNS to manifolds that involve products of spheres, such as those for both boundary point distribution models and s-reps, using the idea of Composite Principal Nested Spheres (CPNS). The idea here is to first develop the PNS representation (Euclideanization) for each spherical component and then concatenate these, together with Euclidean components, into a large Euclidean representation. This would then be followed by PCA on the result.

In the case of s-reps, the spheres in question are one S^2 for each spoke direction and a high-dimensional sphere for the scaled, centered skeletal points. The Euclidean variables are the log spoke lengths and the log scale of the object. The sphere for the n scaled skeletal points, of dimension $3n-4$, arises from centering each skeleton on its center of mass and then scaling each point \mathbf{x}_i by the object scale $\gamma = (\sum_{i=1}^n |\mathbf{x}_i|^2)^{1/2}$.

Of importance is the commensuration between the components before applying the PCA. [Hong, Vicory, Schulz, Styner, Marron & Pizer 2016] designed an experiment to determine the most reasonable commensuration on the Euclideanized features, more precisely, on the object scale and the sphere-resident features derived from the centered, normalized skeletal points. First, the features were transformed to be in the same units. In our case the log-transformed version of the scale factor γ (i.e., $\bar{\gamma} \log(\gamma/\bar{\gamma})$) had units of millimeters, On the other hand, radians (unitless values θ_i for each of the dimensions of the unit sphere) are the units for the Euclideanized shape features derived by PNS from the 2-spheres on which the spoke directions live and from the high-dimensional unit sphere on which the scaled skeletal point tuple for each case live. Thus we multiplied each PNS-derived feature by $\bar{\gamma}$ to put them into units of distance. The problem then is to determine the factor to commensurate the feature capturing scale with the PNS-derived features. We determined that factor by creating a new population that would have a non-varying shape consistent with those in the original population and a scale variation that was the same as those in the original population. To do this, we formed the new population by applying the measured log-transformed γ values for each case to the object of median scale that was then scaled to have $\gamma=1$. By comparing the total variances of the original and created populations, respectively, one could determine the correct commensuration factor between $\bar{\gamma} \log(\gamma/\bar{\gamma})$ and the Euclideanized features from

PNS, namely, $\bar{\gamma}\theta_i$. The experiment concluded that the correct commensuration factor was 1.0, up to sample variation. This idea of separately treating scale can be used for problems of commensuration of other types of variation. In particular, the geometric mean \bar{r}_i of the i^{th} spoke's lengths is used to commensurate PNS-derived features from its spoke directions.

d. Polysphere PCA

Data spaces that are products of spheres, such as the skeletal model spaces of Section 6.1c, have been called polyspheres by [Eltzner, Jung, and Huckemann 2015]. That paper goes on to propose a new method, *Polysphere PCA*. This allows a Euclideanization provided by a more flexible modeling of the dependence between features than that achieved by CPNS's PCA on the geometric properties Euclideanized one by one. This potential improvement is achieved through a distortion of carefully selected angles on each sphere component to map the polysphere data onto a high dimensional sphere and then using PNS on that. This has the advantage relative to CPNS that the PNS is done on a single sphere rather than sphere by sphere, which potentially better captures nonlinear dependence. It has the disadvantage that except near the mean the mapping to the sphere rather severely distorts distances on the manifold. On one population of skeletal shape data this approach has been shown to give a lower dimensional probability distribution representation than is available from CPNS. However, our initial unpublished results comparing classification via Polysphere Euclideanization vs. PNS sphere by sphere Euclideanization [Sharma 2018] suggest that Polysphere Euclideanization's difficulties in separate commensuration for each geometric property can counteract its potential advantages in improved Euclideanization. This holds even when the radius of the various spheres is taken as the geometric means of their entity sizes (length for spokes, scale for skeletal point sequences).

6.2 Classification

A major analytic task when working with populations of shape data is statistical *classification* (also called *discrimination*). For this task, training data, with known class labels, is given, and is used to develop a classification rule for assigning new data to one of the classes. For Euclidean data objects, there are many methods available; see [Duda, Hart & Stork 2001] for a good overview.

The most common methods for classification are based on Euclidean spaces. Particularly widely used these days is the method called support vector machines (SVM); see, for example, [Schölkopf & Smola 2002] for detailed discussion. SVM is based on optimizing the gap in feature space between the training cases in the two classes. A more statistically efficient method called Distance-Weighted Discrimination (DWD), still in a Euclidean feature space, was developed by [Marron, Todd & Ahn 2007]. Its efficiency derives from using all of the training data, not just those near the gap. Both SVM and DWD yield a direction in feature space that optimally separates the classes (Figure 11). Classification

then involves projecting the feature tuple onto the separation direction and deriving the class or the class probability from the resulting scalar value.

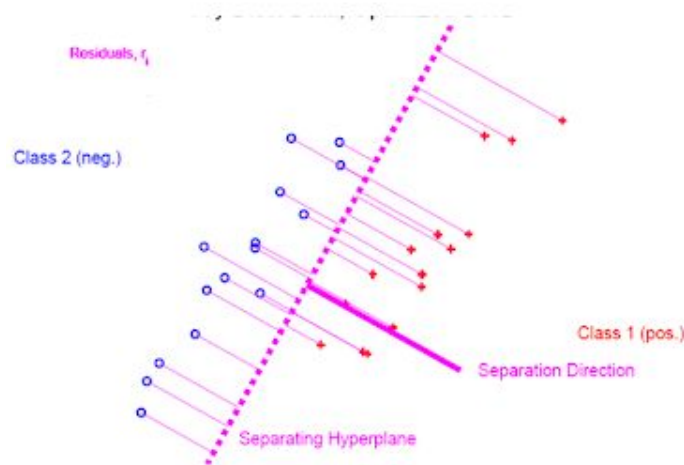


Figure 11. The idea of the separation direction determined from SVM or DWD (each method will determine a somewhat different separation direction)

Use of DWD with data objects lying in a manifold is often done by Euclideanizing the object features and then applying DWD to the result. We call this *Discriminative PNS (DPNS)*. As shown in [Hong, Vicory, Schulz, Styner, Marron, Pizer 2016], DPNS can improve classification accuracy over using the object features directly in DWD.

Section 7 lays out the details of the applications in which the benefits of s-reps statistics have been shown. In brief, diagnostic classifications of two brain structures in two diseases have been shown to be improved by s-reps over boundary point representations, high quality segmentations by posterior optimization of multiple organs in both CT and 3D ultrasound have been produced using shape priors based on s-reps, and useful hypothesis testing by locality and by s-rep feature has been demonstrated.

7. How to compare representations and statistical methods⁵

Given all of the different shape representations discussed in Section 2, the analytic approaches discussed in sections 3 and 4, and the statistical methods discussed in Section 5, an important issue is the comparison of them. One basis for this is knowledge of how they work, which can lead to sensible choices in a wide variety of shape. But it is also interesting to consider quantitative bases for comparison, which is done in this section.

⁵ Much of this Section appeared previously in [Pizer & Marron, 2017]; reprinted here by permission of the copyrighter.

a. Classification accuracy

When a data set is naturally grouped into 2 (or perhaps more) subsets, e.g., pathology vs. normal controls, various methods can be compared on the basis of *classification accuracy*. The classification problem starts with a group of labeled data called the *training set*, and the goal is to develop a rule for classifying new observations. Classification accuracy is simply the rate of correct classifications, either for an independent test set, or using some variation of the cross-validation idea. See [Duda, Hart & Stork 2012] for access to the large literature on classification.

b. Hypothesis testing power

A related approach to comparing shape methodologies, again based on two well-labeled subgroups in the data, is to construct a formal hypothesis test for the difference between the groups. Quantification of the difference then follows from the level of statistical significance, allowing a different type of comparison of the relative merits of various approaches to the analysis. This was done to good effect in [Schulz, Pizer, Marron & Godtliebsen 2015]. The most important contribution there was a technique for dealing in a non-pessimistic way with the high positive correlation that always occurs among geometric properties at nearby points on an object. The methods shown in that paper involve commensurating the various local geometric properties by transforming them to standard Gaussian probabilities. The methods not only work for a test on significance on the whole shape description but also for correcting for multiple tests when testing significance geometric property by geometric property.

c. Specificity, generalization, and compactness

In object statistics work in medicine, two common measures of the quality of a probability distribution derived from data via its eigenmodes are specificity and generalization [Davies 2002a].

Specificity is a measure of how well the estimated probability distribution represents only valid instances of the object. It is computed as the average distance between sample objects randomly generated in the computed shape space and their nearest members of the data.

Generalization is a measure of how close new instances of the object are to the probability distribution estimated from the training cases that do not include the new instance. It is calculated by computing an average over the training cases of the distance of that training case to the shape space specified by eigenmodes, where those eigenmodes are produced using all but the left out training case. That distance is computed as between the left out shape and its projection onto that shape space.

Both specificity and generalization are measured in units of geometric object property differences, where example geometric properties are boundary locations, boundary normal directions, or skeletal spoke directions. In our use of the specificity and generalization measures for s-rep based probability distributions, reported in section 8d, the distances only use corresponding boundary locations and are not based on the

additional features of the full s-rep. In particular, for specificity the measure is the boundary-location-based distance between a random object from the shape space over s-reps to the training object that is nearest according to the boundary location-based distance. For specificity the measure is the boundary-location-based distance between a training object and its projection onto the shape space over s-reps.

Compactness is a measure that quantifies how tight a probability distribution is. Two measures of compactness that are often used are the entropy of the distribution and the determinant of its covariance matrix (total variance).

d. Compression into few modes of variation

The Euclidean PCA decomposition of data into modes of variation is also frequently understood from a signal processing viewpoint. In that case a signal that is actually generated from few dimensions is modified by noise. Noise is high dimensional by definition of having energy across the spectrum. PCA provides a data-driven basis (in the sense of linear algebra) that puts as much of the low dimensional signal as possible into a few basis elements with largest variance. Many of the gains in statistical efficiency, such as those discussed in Section 4, can be understood as providing better signal compression in these terms.

However, this analogy fails in the presence of shape signals, which tend to be widely spread across the PCA spectrum in Euclidean data. In particular, the standard Euclidean assumption of all of the signal being present on in the first few eigenvalues is usually misleading because nonlinear shape signals tend not to lie within a Euclidean space of low dimension so both signal and noise typically spread power among many eigenvalues. The result is that very noisy data can be measured to require fewer eigenmodes to achieve a given fraction of total variance than less noisy data. Hence standard *dimension reduction* approaches, based on “total signal power” or “percent of variation explained”, are usually inappropriate.

Yet there is still a natural desire to think in terms of *effective dimensionality*, i.e., the concept that the true underlying signal has much less variation than is present in noisy data. Ideas based in *random matrix theory* are promising avenues for research in this direction. See [Yao, Bai & Zheng 2015] for good discussion of using this powerful theory in the context of Euclidean PCA. The first important part of this theory is the asymptotic probability distribution of the full collection of eigenvalues under a pure noise model, called the Marčenko-Pastur distribution [Marčenko & Pastur 1967]. The second part is the corresponding limiting distribution of the largest eigenvalue, called the Tracy Widom distribution [Tracy & Widom 1996].

As yet, little development of Marčenko-Pastur theory for data on curved manifolds appears to have been done. An exception is unpublished research in progress by Choi, Pizer and Marron. That work is studying analogs of limiting distributions of eigenvalues in the case of PNS. Eigenvalues are defined in terms of sums of squares of scores, i.e., of the energy measured along the manifold surface. Careful study has revealed appropriate

scale factor modifications that take the curvature of the space appropriately into account. These yield variations of the Marčenko-Pastur distribution that, as illustrated in Figure 12, are seen to effectively fit simulated data on the high dimensional sphere very well. Both panels are Q-Q plots showing quantiles of the theoretical Marčenko-Pastur distribution on the vertical axis as a function of quantiles of the empirical distributions on the horizontal axis. A good distributional approximation is indicated by the points lying close to the red 45° line.

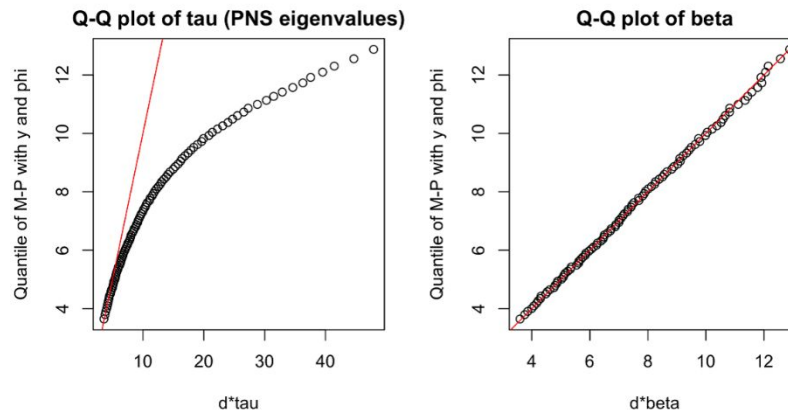


Figure 12. Q-Q plots showing poor goodness of fit of the Marčenko-Pastur distribution to the raw PNS eigenvalues in the left panel and good fit to the rescaled eigenvalues in the right panel. The data consists of simulated Gaussian noise on a $d=1000$ -dimensional tangent plane projected to the unit sphere with sample size $n=100$. The notation “tau” refers to PNS eigenvalues, and the notation “beta” refers to the rescaled version.

8. Results of classification, hypothesis testing, and probability distribution estimation⁶

A good approach to comparing shape analysis methods is to study their impact when used for various applications. This Section reviews some recent work, with many specific applications of the above ideas.

a. Classification of mental illness via hippocampus and caudate s-reps

[Hong, Vicory, Schulz, Styner, Marron & Pizer 2016] compared boundary point distribution models (PDMs) vs. s-reps and Euclideanization vs. direct Euclidean analysis of the ambient space coordinate values in classifying a hippocampus as to whether it was from a typical individual (control) or from a first-episode schizophrenic. The hippocampus was represented using a 3×8 skeletal grid (on each side of the skeleton), and the analysis was done using 238 schizophrenics and 56 controls. They showed that, according to areas under the Receiver Operating Characteristic curve (ROC) [Hanley & McNeil 1982], Euclideanizing boundary PDMs produced better classification than

⁶ Much of parts a-d of this Section appeared previously in [Pizer & Marron, 2017]; reprinted here by permission of the copyrighter.

without Euclideanization and that Euclideanized s-reps produced better classification than either of the boundary-based analyses (Figure 13 and Table 8.1). In Table 8.1, as well as Tables 8.2-3, “AUC” stands for the area under the ROC (the measure of classification quality), “PNS” stands for Euclideanization by Principal Nested Spheres, and “DWD” stands for the Euclidean classification method used, Distance-Weighted Discrimination.




Figure 13. ROCs for classifying hippocampi as to schizophrenia vs. typical using s-reps and PDMs, each with original features and their Euclideanized counterparts.

Methods	Hippocampus AUC	Caudate AUC
s-reps + PNS + DWD	.72	.75
s-reps + DWD	.67	---
boundary s-rep PDMs + PNS + DWD	.70	.72
boundary s-rep PDMs + DWD	.67	---
Random guessing	.5000	.5000

Table 8.1. AUCs for classifying hippocampi or caudate nuclei between individuals having experienced a first schizophrenic behavioral episode and controls. The classifications compared are based on s-reps, respectively for row 1) s-reps analyzed using PNS Euclideanization, 2) s-reps analyzed without Euclideanization, 3) spoke end boundary points analyzed using PNS Euclideanization, 4) spoke end boundary points analyzed without Euclideanization, 5) pure random guessing

As also shown in Table 8.1, similar results were obtained on the same population of subjects for the caudate nucleus that was represented using 4×8 grid of skeletal points.

Hong et al. also showed the usefulness of displaying the variation of the object along the vector in the Euclideanized feature space passing through the pooled mean and in the separation direction (see Figure 11) of the classes (Figure 14).




Figure 14. Boundaries implied by s-reps of hippocampi: Left: at pooled mean - 2 standard deviations in the separation direction for classification between first episode schizophrenics and controls. Right: at pooled mean + 2 standard deviations in that separation direction. The shape change between these can be particularly noted by the changed bending in the right third of the hippocampus.

Hong et al. [2018] further evaluated s-reps for the classification of autism vs. controls via hippocampal or caudate shape imaged via MRI for at-risk subjects at the age of 6 months. That is, the classification was between babies who would show autistic behavioral symptoms two years later and those who would not. The hippocampus was represented using a 3×8 skeletal grid (on each side of the skeleton) and the caudate nucleus was represented using a 4×8 skeletal grid (on each side of the skeleton). The analysis was done using 49 subjects who eventually displayed autistic symptoms and 149 who did not display such symptoms.

Methods	Hippocampus AUC	Caudate AUC
s-reps + PNS + DWD	.6400	.5708
s-reps + DWD	.6123	.5419
boundary s-rep PDMs + PNS + DWD	.6062	.5400
boundary s-rep PDMs + DWD	.6050	.5365
Global volume + DWD	.5560	.5372
Random guessing	.5000	.5000

Table 8.2. AUCs for classifying hippocampi or caudate nuclei between children at high risk for Autism Spectrum Disorder actually developing autism symptoms and those not developing autism symptoms. The classifications compared are based on s-reps, respectively for row 1) s-reps analyzed using PNS Euclideanization, 2) s-reps analyzed without Euclideanization, 3) spoke end boundary points analyzed using PNS Euclideanization, 4) spoke end boundary points analyzed without Euclideanization, 5) s-rep-implied volume only and 6) pure random guessing

Similarly to Table 8.1, Table 8.2 shows that 1) s-reps analyzed with PNS yield notably better classification than with boundary points; 2) much less gain in classification quality occurs when the s-reps are not Euclideanized; 3) little gain in classification quality occurs when the boundary points are not Euclideanized; 4) at least for the hippocampus, global volume gives inferior classification relative to any of the shape representations.

Hong's results above were on unrefined s-reps. In yet unpublished work Wang [2019] has shown that when the analysis is repeated on those s-reps after refinement, the same AUC value is obtained. That is, the measure of the classification neither increases nor decreases.

Finally, [Hong 2018] showed how to better skeletally represent an object like the caudate nucleus that approximately ends at a cusp of zero width. He developed a new primitive that represented one end of the s-rep narrowing to a cusp. That primitive consists simply of the location of the cusp and the tangent direction of the cusp. He showed how to interpolate a skeleton of an object with the grid s-rep augmented by the new primitive and how to fit such a skeleton to caudate nuclei. When these new skeletons were used for the caudate classification in the autism data, the AUC (in the upper-right entry of Table 8.2) became .587; that is, it added .016 to the AUC derived from the s-rep without the cusp primitive.

For all of these area under the ROC measurements one would wish error bars. Unfortunately, it is not known how to compute error bars for this measurement when it is made with repeated holdouts, which results in complicated correlations among the values for different holdouts. Our group is researching what other measurements might allow the computation of error bars. At present, all we can say is that consistently across two anatomic objects and two diseases the addition of the object width and direction features provided by s-rep spokes increases the area under the ROC in classifications by a few hundredth values. Lest one feel this is a small difference, note that it is around $\frac{3}{4}$ of the difference in AUC between random guessing and classification using global volume of the object.

Looking at the classification results together, we see that for boundary point representation, PNS Euclideanization makes only a small improvement in AUC, but for s-reps it makes a notable improvement in AUC. And using s-reps to capture not only location information but also object widths and spoke directions consistently yields notable improvement in AUC over using locational information alone.

b. Hypothesis testing via s-reps

[Schulz, Pizer, Marron & Godtlielsen 2015] demonstrated the ability to test hypotheses on shape variation between two objects using s-rep features. They showed how to do not only global hypothesis tests but also GOP (geometric object property)-by-GOP and location-by-location tests. The method involved a modification of the statistical technique of permutation tests in a way that recognizes that means benefit from being backward means and that GOP differences need to use a metric appropriate for curved manifolds. As discussed in Section 7b, this work reported a new method for compensating for the notable correlations of these GOPs. With this approach they were able to analyze which locations on hippocampi of first-episode schizophrenics had statistically significant GOP differences from controls, and which GOPs had those statistically significant differences. Thus they found important shape differences between the classes.

c. Shape change statistics

[Vicory 2016] studied statistics on the change of object shape between two stages. Realizing that statistics on change requires transporting each object pair such that the starting object was at a fixed place on the manifold of objects, he solved this problem by pooling all $2n$ objects in the n shape pairs in the training set, producing a polar system by PNS on the pool of GOP tuples, Euclideanizing each object according to that polar system, and then producing Euclidean (ordinary) differences between the Euclideanized features of each pair.

Vicory [2016] studied this method in two applications. The first was on ellipsoids that were randomly bent and/or twisted. The GOPs in this study were the boundary point coordinates provided by spherical harmonic analysis. His findings with simulated data produced from random samples from a single transformation type were as follows: 1) Analysis either with or without Euclideanization yielded a single correct eigenmode. 2) Analysis with Euclideanization allowed more accurate estimation of an object's bending or twisting angle. 3) When the mode of variation was visualized, the object moved more naturally than using PCA alone. He also found that when the simulated data came from random bending cascaded with twisting, both types of analysis yielded two modes of variation but Euclideanization allowed a more correct estimate of the relative variances and mean deformations that were closer to the expected deformations than their PCA-alone counterparts.

Vicory's second test was on prostates segmented from MRI and the same prostate within 3D transrectal ultrasound; the latter was deformed by the ultrasound transducer (discussed in more detail in Section 8e). The shape-change eigenmodes were used as a shape space within which to segment the prostate from the ultrasound image, given the patient's prostate shape derived from his MRI. The geometric object properties he used were the skeletal points, spoke lengths, and spoke directions from fitted s-reps. He found that using the shape-change space resulting from Euclideanization followed by subtraction yielded more accurate segmentations than when a shape space was formed by s-rep feature differences applied to the mean of the prostates in the training MRIs and then applying CPNS to the resulting objects. For a target segmentation using that superior method, it was necessary to form an ultrasound-relevant starting prostate s-rep from which the optimization over the shape-change shape space would be computed. That initialized s-rep in the ultrasound image was computed by starting with the patient's in-MRI prostate s-rep and then applying to it (adding in the Euclideanized form) the mean prostate s-rep difference between the ultrasound-based prostates and their corresponding MRI prostates.

Recently Hong has applied this shape-change Euclideanization approach to two classes of pairs of shapes, at two different ages, namely 6 and 12 months. For each anatomic object he analyzed, he pooled all his objects for children at both 6 months and 12 months to yield a Euclideanization. Then he analyzed shape changes via this Euclideanization to determine how informative shape differences in the hippocampus and caudate nucleus,

respectively were between children at high risk for autism who do not develop autistic symptoms and those at high risk who do develop such symptoms. However, as shown in Table 8.3 (from [Hong 2018]), the classification based on the shape difference provides poorer classification than either of the classifications of 6-month shapes and 12-month shapes. As previously noted, in that table “AUC” stands for the area under the ROC (the measure of classification quality), “PNS” stands for Euclideanization by Principal Nested Spheres, and “DWD” stands for the Euclidean classification method used, Distance-Weighted Discrimination.

Methods	Hippocampus change AUC	Caudate change AUC
temporal difference s-reps + PNS + DWD	.5427	.5041
6 month s-reps + PNS + DWD	.6400	.5087
12 month s-reps + PNS + DWD	.5837	.5605
random guessing	.5000	.5000

Table 8.3. AUCs for classifying between children at high risk for Autism Spectrum Disorder actually developing autism symptoms and those not developing autism symptoms. The classifications compared are based on hippocampi or caudate nucleus s-reps, respectively for 1) temporal shape differences at 6 months and 12 months, 2) 6 months’ shapes, 3) 12 months’ shapes, and 4) pure random guessing

d. Correspondence evaluation via entropy minimization

The basic idea of correspondence optimization is to consider all reparametrizations of each object in a population and to find the reparametrization of each that produces the tightest estimated probability distribution. The notion of reparametrizing an object representation is to apply a diffeomorphism on the boundary surface or skeletal surface while leaving the object geometrically the same. For a spatially sampled object boundary, that is a fancy way of saying how to shift the respective sample points along the boundary. For example, if an object of spherical topology is understood as a smooth 1-to-1 mapping from the unit sphere’s boundary to the object boundary, the shifting of where each sphere point maps to on the object is a reparametrization [Kurtek, Srivastava et al. 2013, Brechbueler, Gerig et al., 1995]. For s-reps the shifting is of spokes along the skeletal surface.

The tightness of a probability distribution can be measured in many ways, but the ones that have turned out to be the most effective are based on information theory. Taylor and his team [Davies 2002b] pioneered a form based on minimum description length (MDL), and an almost equivalent form based on entropy was developed in Whitaker’s laboratory [Cates 2007]. In these works one uses objects represented according to the specified parameterization, and one fits the population of GOPs derived from these parameterizations by a Gaussian distribution. Then the entropy of this Gaussian is minimized over all object reparameterizations, while also penalizing irregularity of the discrete object samples (points on a boundary or spokes on a skeleton).

The entropy of an n -dimensional Gaussian is given by

$$n \left(\frac{1}{2} \ln(2\pi) + \text{the mean of the logarithms of the principal variances} \right).$$

A difficulty is that when the idea is applied to a probability distribution estimated from data by a PCA or PCA-like method that involves Euclideanization, the successive sorted principal variances get successively smaller and eventually are dominated by noise in the data. Thus, each logarithm of these small principal variances is very large in magnitude, so these noise-induced principal variances dominate the population principal variances in the entropy calculation. Cates et al. [2007] dealt with this problem by adding a small constant to each principal variance. [Tu, Styner, Vicory, Paniagua, Prieto, Yang, Pizer 2015] dealt with it by cutting off the series when the ratio of the principal variance to the total variance fell below a certain threshold.

Cates et al. [2007] showed that entropy minimization for GOPs that were a boundary point distribution could yield improvements in hypothesis testing significances. In their work boundary points were slid along the boundary to minimize entropy. [Tu, Styner, Vicory, Paniagua, Prieto, Yang, Pizer 2015] developed the same idea for s-reps, whereby the skeletal spokes were slid along the skeletal surface to minimize entropy. They showed (Figure 15) the result that, according to the measures described in Section 5c on a training set of hippocampi, when the s-rep-based correspondence method was applied, the boundary points at the ends of the spokes had better (lower) values of the statistical properties: specificity and generalization (see Section 7c) than when the Cates boundary point shifting method was applied. This result is surprising and arguably indicative of the strength of s-reps because the object comparisons used in measuring specificity and generalization were based on boundary point distributions, so one might have expected that optimization directly of boundary point shiftings, i.e., the Cates method, would be superior.

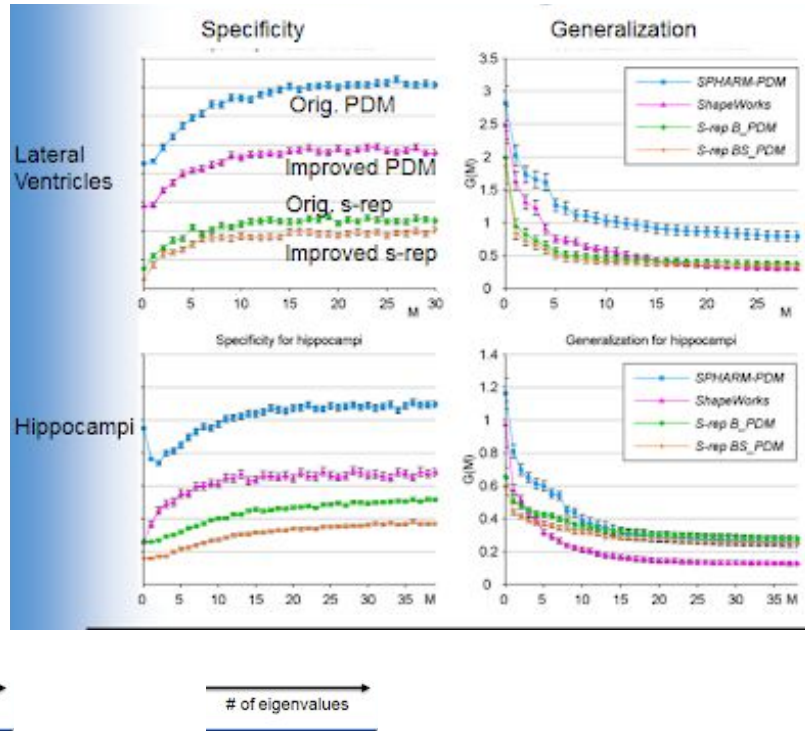


Figure 15. Generalization and specificity improvements due to s-rep-based and PDM-based correspondence optimization. Smaller values indicate better performance.

e. Segmentations by posterior optimization

An important application of shape statistics is *segmentation* in medical images, where the goal is to find the region of an image occupied by an object such as a particular organ. A series of successive improvements in shape analysis resulting in improved image segmentation can be found in [Pizer 2001, 2013], [Fletcher 2004], [Rao 2005], [Gorcowski 2007], and [Vicory 2014, 2016].

For many years Pizer, Chaney, et al. (Chapter 9 in [Siddiqi and Pizer 2008]) worked on the segmentation of male pelvis organs from CT to aid radiation treatment planning for prostate cancer. Their method involved an earlier form of s-reps and used them in both terms of a posterior optimization approach, slightly modified to make the log prior and log likelihood terms have the same variance. In the log prior (shape statistics) term each organ's s-rep statistics was used, and in the log likelihood (statistics of the match between the image and the object geometry), intensity-derived features at positions implied by interpolated s-rep spokes were used to produce statistics by PCA. This method was taken over by a spinoff company, Morphormics, now part of Accuray, Inc. With further modifications of the method by Morphormics, the evaluations of the method showed it to be superior to various competitors [Morphormics 2008] and were brought into broad clinical use.

In another segmentation application, Vicory [2016] worked on extracting the prostate from 3D trans-rectal ultrasound (TRUS), in which the intra-rectal probe deforms the adjacent prostate. The method was given the shape of that patient’s prostate in an MR image that did not have the ultrasound probe in place. His term measuring the match between the image and the object geometry in the objective function he optimized that involved intensity and as well, 48 texture features. It classified each pixel according to that 49-tuple, with the result being the probability, $p(\text{inside})$, that the pixel was inside the prostate (see the left panel of Figure 16). Then his term measuring the match between the image and the object geometry for paired regions respectively inside and outside a particular patch within the s-rep-implied boundary was

$$\sum_{\text{patches } i} \left[\left(\rho_i^{\text{in}} - \rho_i^{\text{out}} \right) - \left| \rho_i^{\text{in}} + \rho_i^{\text{out}} - 1 \right| \right],$$

where ρ_i^{in} and ρ_i^{out} are the average $p(\text{inside})$ values for the regions inside and outside, respectively, the i^{th} patch. Each patch was a quadrilateral centered at an s-rep spoke tip and extending halfway to the adjacent spoke in each of the two cardinal directions of the skeletal surface’s grid of points. Thus, it used the correspondence across prostates provided by the s-rep.

Vicory’s log prior used the shape change statistics method described in Section 8c, namely on the change between the prostate s-rep in the patient’s MRI and that in the patient’s 3D TRUS image. It was intended that this method be followed by a local refinement step, but this refinement was never studied. Nevertheless, even without refinement the method’s segmentations (see Figure 16) were competitive with others that had been proposed (e.g., [Garnier 2011], [Qiu 2014]).




Figure 16. Left: an example of the $p(\text{inside})$ image of a prostate derived from a 3D TRUS image; Middle: the prostate before the shape change; Right: the resulting prostate segmentation from the 3D TRUS image, i.e., after the shape change.

9. The code and its performance

The code for fitting an s-rep with a skeletal surface to an object boundary, s-rep visualization, and s-rep statistics (via PNS) is becoming part of Slicer Shape AnaLysis Toolbox (SlicerSALT), a web-resident toolkit that was first announced at MICCAI 2018. While the previous versions of code for these purposes were hard to understand and use, the version in the toolkit is designed for clarity of use and modification. Moreover, the computer representation of the s-rep used in that toolkit is more generalizable and closer to the mathematics than the “legacy representation” used in the previous versions of the code⁷. In the SlicerSALT s-rep code [Liu 2019] automatic fitting of an s-rep to an object boundary takes approximately 5 to 10 minutes. In that code the input object is given as a triangle mesh, but extension to representations of the input object in the form of a binary image or in the form of a signed distance-to-boundary image are anticipated.

The generalizations of s-reps that will ultimately be included in the Slicer-SALT package will include generalized cylinders, i.e., quasi-tubular objects that are well represented by a skeletal curve rather than a surface.

The s-rep code in SlicerSALT is for single objects, but we have ongoing research on multi-object models, i.e., models that don’t simply concatenate the individual objects’ s-reps but also capture s-rep-derived geometry on the relations between the objects [Damon 2017]. Because objects can slide along each other, this is a challenging problem, yet one we are making progress on.

Some objects have essential branching of the skeletal surface. The branches represent essential protrusions or essential indentations. Because so many anatomic objects can be represented by s-reps without branching, there has been little work on those with essential branching. What little work on s-reps for objects with protrusions or indentations has treated each branch as having its own s-rep and has provided a way for the branch to be additively (for a protrusion) or subtractively (for an indentation) connected to the parent s-rep [Han 2008]. Yushkevich and his colleagues have worked on branching with their cm-reps [Pouch 2015].

1. Weaknesses of the skeletal approach

While s-reps have been shown to be a powerful representation for applications using shape statistics, some objects for which statistics are needed are not suitably represented by a skeletal model. These include objects that need variable skeletal topology, for example, that vary with some instances of a skeletal surface becoming so narrow that they approach a curve. They also presently include objects with variation in the branching arrangement, for which tree statistics would need to be combined with that on the skeletal descriptions of the branches; methods of tree statistics, while of great interest, are still

⁷ The s-reps displayed in figs. 1 and 6 are of the form used in the toolkit, but the s-rep-implied boundaries shown in a variety of figures were all created using the legacy form of the s-rep.

under development [Aydin et al. 2009, Shen et al. 2014, Skwerer et al. 2014, Bendich et al. 2016].

Open questions on s-reps include

- 1) quantitative comparisons between Yushkevich’s cm-reps and s-reps;
- 2) the production of statistical methods on s-reps that work directly on the polysphere manifold;
- 3) the use of other statistical analysis techniques such as Fletcher’s Bayes [Zhang 2013] or Sommer’s Brownian motion induced Gaussians [Sommer 2015] in s-rep applications.
- 4) improvements in the representation of generalized cylinders (“quasi-tubes”) via s-reps and doing statistics thereon.

Acknowledgements

We are grateful to Martin Styner for useful conversations and for providing brain MRI data, IBIS for providing brain MRI data, and the editors of this book for helpful suggestions in the writing.

References

- Aydin, B, G. Pataki, H Wang, E Bullitt, JS Marron (2009). A principal component analysis for trees. *The Annals of Applied Statistics*, 1597-1615.
- Aylward, S, E. Bullitt (2002). "Initialization, noise, singularities, and scale in height ridge traversal for tubular object centerline extraction" *IEEE Transactions on Medical Imaging* 21(2):61-75.
- Bendich, P, JS Marron, E Miller, A Pieloch, S Skwerer (2016). Persistent homology analysis of brain artery trees. *The annals of applied statistics*, 10(1), 198.
- Blum, H (1967). A transformation for extracting new descriptors of shape. In W. Walthen-Dunn, ed., *Models for the Perception of Speech and Visual Form*. Cambridge MA, MIT Press.
- Brechtbuehler, B, G Gerig, O Kuebler (1995). Parametrization of closed surfaces for 3-d shape description. *Computer Vision and Image Understanding*, 61(2):154–170.
- Cates, J, PT Fletcher, M Styner, M Shenton, R Whitaker (2007). “Shape modeling and analysis with entropy-based particle systems,” in *Information Processing in Medical Imaging*, pp. 333–345, Springer.
- Damon, J (1999). “Properties of Ridges and Cores for Two-Dimensional Images” *Jour. Math. Imaging and Vision* 10, pp 163–174

Damon, J, & Marron, J. S. (2014). Backwards principal component analysis and principal nested relations. *Journal of Mathematical Imaging and Vision*, 50(1-2), 107-114.

Damon, J, E Gasparovic (2017). Modeling multi-object configurations via medial/skeletal linking structures. *Int. J. Comp. Vis.* 124, pp 255-272.

Davies, R (2002a). *Learning shape: optimal models for analysing natural variability*. Ph.D. dissertation, Dept. of Imag. Sci. and Biomed. Eng., Univ. of Manchester, England, 2002.

Davies, RH CJ Twining, TF Cootes, JC Waterton, CJ Taylor (2002b), "A minimum description length approach to statistical shape modeling," *IEEE Trans. Med. Imag.*, vol. 21, no. 5, pp. 525-537.

Duda, R. O., Hart, P. E., & Stork, D. G. (2012). *Pattern classification*. John Wiley & Sons.

Eltzner, B., Jung, S., & Huckemann, S. (2015, October). Dimension Reduction on Polyspheres with Application to Skeletal Representations. In *International Conference on Networked Geometric Science of Information* (pp. 22-29). Springer International Publishing.

Fletcher, PT (2004). *Statistical Variability in Nonlinear Spaces: Application to Shape Analysis and DT-MRI*, PhD Dissertation, Dept. of Computer Science, University of North Carolina, Chapel Hill.

Fletcher, P. T., Lu, C., Pizer, S. M., & Joshi, S. (2004). Principal geodesic analysis for the study of nonlinear statistics of shape. *Medical Imaging, IEEE Transactions on*, 23(8), 995-1005.

Fletcher, P.T. (2013) Geodesic regression and the theory of least squares on Riemannian manifolds, *IJCV*, 105 (2), pp. 171-185.

Fridman, Y, SM Pizer, S Aylward, and E Bullitt (2004). "Extracting Branching Tubular Object Geometry via Cores," *Medical Image Analysis Journal, Special MICCAI issue*. 8(3): 169-176.

Garnier, C J-J Bellanger, K Wu, H Shu, N Costet, R Mathieu, R De Crevoisier, J-L Coatrieux (2011). "Prostate segmentation in hifu therapy." *IEEE Transactions on Medical Imaging*, vol. 30, no. 3, pp. 792-803.

Gorcowski, K, M Styner, J-Y Jeong, JS Marron, J Piven, Hr Cody Hazlett, SM Pizer, G Gerig (2007). [Statistical Shape Analysis of Multi-Object Complexes](#)," *Proc. CVPR*, June 2007 pp. 1-8.

- Han Q (2008). *Proper Shape Representation of Single Figure and Multi-Figure Anatomical Objects*, Computer Science dissertation, University of North Carolina at Chapel Hill.
- Hanley, JA, BJ McNeil (1982). The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*, 143(1), 29-36.
- Hong, J., Vicory, J., Schulz, J., Styner, M., Marron, J., & Pizer, S. M. (2016). Classification of medically imaged objects via s-rep statistics. *Med. Image Anal.*, 31, 37–45.
- Hong, JP (2018). *Classification of Neuroanatomical Structures Based on Non-Euclidean Geometric Object Properties*, Computer Science dissertation, Univ. of North Carolina at Chapel Hill.
- Jeong, J. Y. (2009). *Estimation of probability distribution on multiple anatomical objects and evaluation of statistical shape models*. Doctoral dissertation, The University of North Carolina at Chapel Hill.
- Jung, S., Dryden, I. L., & Marron, J. S. (2012). Analysis of principal nested spheres. *Biometrika*, 99(3), 551-568.
- Katz, R (2002). *Form metrics for interactive rendering via figural models of perception*. Computer Science dissertation, University of North Carolina at Chapel Hill.
- Kurtek, S, Z. Ding, E. Klassen, and A. Srivastava (2011). Parameterization-Invariant Shape Statistics and Probabilistic Classification of Anatomical Surfaces. *Information Processing in Medical Imaging (IPMI)*, LNCS 10265. Springer.
- Liu, Z., J Vicory, B Paniagua (2019). Shape analysis via skeletal models. SALT. At URL <https://github.com/KitwareMedical/SlicerSkeletalRepresentation/blob/master/SrepModule-Tutorial-Visualizer.pdf>.
- Marčenko, VA, LA Pastur (1967). Distribution of eigenvalues for some sets of random matrices. *Mathematics of the USSR-Sbornik*, 1(4), 457.
- Ma, R, Q Zhao, R Wang, J Damon, J Rosenman, S Pizer (2018). Generalized Cylinder Deformation for Visualization of a Tube Interior. Univ. of North Carolina Dept. of Computer Science Tech. report.
- Marron, J. S., Todd, M. J., & Ahn, J. (2007). Distance-weighted discrimination. *Journal of the American Statistical Association*, 102(480), 1267-1271.

Morphormics (2008). Morphormics Technology Overview White Paper. Available on www.morphormics.com until 2012. Morphormics, Inc. (now Accuray, Inc.)

Morse BS, SM Pizer, DS Fritsch (1994). "Robust Object Representation Through Object-Relevant Use of Scale," *Proceedings SPIE Medical Imaging '94: Image Processing*, **2167**, 104-115.

Patrangenaru, V., & Ellingson, L. (2015). *Nonparametric statistics on manifolds and their applications to object data analysis*. CRC Press.

Pennec, X. (2015, October). Barycentric subspaces and affine spans in manifolds. In *International Conference on Networked Geometric Science of Information* (pp. 12-21). Springer, Cham.

Pizer S, S Joshi, PT Fletcher, M Styner, G Tracton, Z Chen (2001). "Segmentation of Single-Figure Objects by Deformable M-reps," *Medical Image Computing and Computer-Assisted Intervention (MICCAI 2001)*, WJ Niessen, MA Viergever, eds. *Lecture Notes in Computer Science*, **2208**: 862-871, 2001.

Pizer, S.M., Jung, S., Goswami, D., Vicory, J., Zhao, X., Chaudhuri, R., Damon, J.N., Huckemann, S. and Marron, J.S. (2013). Nested sphere statistics of skeletal models. In *Innovations for Shape Analysis* (M. Breuss, A. Bruckstein, and P. Maragos, eds.), pp. 93-115. Springer Berlin Heidelberg.

Pizer, S. M., & Marron, J. S. (2017). Object Statistics on Curved Manifolds. Chapter 13 in G Zheng, S Li, G Székely, *Statistical Shape and Deformation Analysis: Methods, Implementation and Applications*, pp 137-164, Academic Press.

Pouch, AM, S Tian, M Takebe, J Yuan, R Gorman, AT Cheung, H Wang, BM Jackson, JH Gorman, RC Gorman, PA Yushkevich (2015). Medially constrained deformable modeling for segmentation of branching medial structures: application to aortic valve segmentation and morphometry. *Medical Image Analysis*, 26(1), pp.217-231.

Qiu, W, J Yuan, E Ukwatta, Y Sun, M Rajchl, A Fenster (2014) "Prostate segmentation: An efficient convex optimization approach with axial symmetry using 3-d TRUS and MR images," *IEEE Transactions on Medical Imaging*, vol. 33, no. 4, pp. 947-960.

Rao, M, J Stough, Y-Y Chi, K Muller, GS Tracton, SM Pizer, EL Chaney (2005). "Comparison of Human and Automatic Segmentations of Kidneys from CT Images," *International Journal of Radiation Oncology, Biology, Physics*, vol. 61, no. 3, pp. 954-960, 2005.

Schölkopf, B., & Smola, A. J. (2002). *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press.

Schulz, J, SM Pizer, JS Marron, F Godtlielsen (2015). "Nonlinear Hypothesis Testing of Geometrical Object Properties of Shapes Applied to Hippocampi," *J Math. Imag. & Vision*, 54(1), pp 15-34.

Sen, S. K., Foskey, M., Marron, J. S., & Styner, M. A. (2008, May). Support vector machine for data on manifolds: An application to image analysis. In *Biomedical Imaging: From Nano to Macro, 2008. ISBI 2008. 5th IEEE International Symposium on* (pp. 1195-1198). IEEE.

Sharma, A (2018). Euclideanization of s-reps via Discriminative Principal Nested Spheres (DPNS) vs. Polysphere Nested Deformed Spheres (PNDS). Internal Report, Dept. of Computer Science, Univ. of NC.

Shen, D, H Shen, S Bhamidi, Y Muñoz Maldonado, Y Kim, Y, JS Marron (2014). Functional data analysis of tree data objects. *Journal of Computational and Graphical Statistics*, 23(2), 418-438.

Siddiqi, K & SM Pizer (2008), *Medial Representations: Mathematics, Algorithms and Applications*. Springer Publishers, 2008

Skwerer, S, E Bullitt, S Huckemann, E Miller, I Oguz, M Owen, JS Marron (2014). Tree-oriented analysis of brain artery structure. *Journal of Mathematical Imaging and Vision*, 50(1-2), 126-143.

Sommer, S. (2015). Anisotropic Distributions on Manifolds: Template Estimation and Most Probable Paths, Lecture Notes in Computer Science, vol. 9123, S. Ourselin et al., eds., *Information Processing in Medical Imaging*, Springer, 193-204.

Tracy, CA, Widom H (1996). On orthogonal and symplectic matrix ensembles. [*Comm. Math. Phys.*](#) 177, 727–754.

Tu, L, J Vicory, S Elhabian, B Paniagua, JC Prieto, JN Damon, R Whitaker, M Styner, SM Pizer (2016). "[Entropy-based Correspondence Improvement of Interpolated Skeletal Models](#)," *Computer Vision and Image Understanding*, vol. 151, no. , pp. 72-79.

[Vaillant M](#), [MI Miller](#), [L Younes](#), [A Trouvé A](#) (2004). Statistics on diffeomorphisms via tangent space representations. *Neuroimage*, vol 23, pp. S161-S169.

Vicory, J (2014). Xxx

Vicory, J (2016). [Shape Deformation Statistics and Regional Texture-based Appearance Models for Segmentation](#). Computer Science dissertation, University of North Carolina at Chapel Hill.

Yao, J, S Zheng, Z Bai (2015). *Large sample covariance matrices and high-dimensional data analysis*. [Cambridge Series in Statistical and Probabilistic Mathematics, 39](#). Cambridge University Press, New York.

Yushkevich, PA, H Zhang, JC Gee (2006). Continuous medial representation for anatomical structures. *IEEE transactions on medical imaging*, 25(12), pp.1547-1564.

Wang, J, Z Liu, JP Hong, SM Pizer, Hippocampus classification via refined vs. initially fit s-reps (2019). Internal Report, Dept. of Computer Science, Univ. of NC.

Zhang M and P. Thomas Fletcher, Probabilistic Principal Geodesic Analysis, in Proceedings of Advances in Neural Information Processing Systems 26, pp. 1178-1186, 2013

Zhao, Q, T Price, S Pizer, M Niethammer, R Alterovitz, J Rosenman (2015). "[Surface Registration in the Presence of Missing Patches and Topology Change](#)," in *Medical Image Understanding and Analysis Conference*.