



Published in final edited form as:

Comput Vis Image Underst. 2016 October ; 151: 72–79. doi:10.1016/j.cviu.2015.11.002.

Entropy-based Correspondence Improvement of Interpolated Skeletal Models

Liyun Tu^{a,b,1}, Jared Vicory^{b,1,*}, Shireen Elhabian^c, Beatriz Paniagua^b, Juan Carlos Prieto^d, James N. Damon^b, Ross Whitaker^c, Martin Styner^b, Stephen M. Pizer^b

^aChongqing University, Shapingba, China

^bUniversity of North Carolina at Chapel Hill, Chapel Hill, NC, USA

^cUniversity of Utah, Salt Lake City, UT, USA

^dBrigham and Women's Hospital, Boston, MA, USA

Abstract

Statistical analysis of shape representations relies on having good correspondence across a population. Improving correspondence yields improved statistics. Point distribution models (PDMs) are often used to represent object boundaries. Skeletal representations (s-reps) model object widths and boundary directions as well as boundary positions, so they should yield better correspondence.

We present two methods: one for continuously interpolating a discretely-sampled skeletal model and one for improving correspondence by using this interpolation to shift skeletal samples to new positions. The interpolation operates by an extension of the mathematics of medial structures. As with Cates' boundary-based method, we evaluate correspondence in terms of regularity and shape-feature population entropies.

Evaluation on both synthetic and real data shows that our method both improves correspondence of s-rep models fit to segmented lateral ventricles and that the combined boundary-and-skeletal PDMs implied by these optimized s-reps have better correspondence than optimized boundary PDMs.

Keywords

Statistical Shape Analysis; Correspondence; Skeletal Models

*Corresponding author. Phone: +1 919 886 5757 jvicory@cs.unc.edu.

¹These authors contributed equally to this work.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

1. Introduction

Statistical analysis of shape populations is an important task in many image analysis applications. Achieving good correspondence [1] across a population of shapes is a necessary step in computing accurate statistics. For point distribution models (PDMs) [2], this means having the points evenly spread around the object while each point has similar local geometry on every object. The most common use of the PDM is to represent the boundary of an object. Cates et al. [3] developed a method for improving correspondence of a set of PDMs by shifting points along the boundary to minimize geometric entropy less a regularity entropy summed over the population, which is implemented in the publicly available *ShapeWorks* [4] software. This optimization has been shown to produce improved statistics on populations of PDMs.

In many applications [5, 6, 7], it is beneficial to model object interiors as well as boundaries. Skeletal representations (s-reps) [8] are one method for doing this. An s-rep consists of a discretely sampled skeletal surface with vectors called spokes pointing from the skeleton to the object boundary. As with PDMs, good correspondence is required for statistical analysis of s-reps. At the same time, it would be interesting to compare the statistical properties of a boundary PDM put into correspondence directly by the Cates method to one implied by the s-reps correspondence. In this paper, we present a method for improving correspondence of s-reps by shifting spokes along the skeleton. Similarly to the boundary-based method, we shift spokes to minimize a geometric entropy less a regularity entropy summed over the population. In order to shift these discrete spokes, we develop a method for interpolating a continuous s-rep from the discrete samples. The interpolation is a new result generalizing the mathematics of medial structures [9].

We evaluate our method on a set of 31 lateral cerebral ventricles. We show that our method improves the correspondence of s-reps fit to this data. We also compare the results of the s-rep correspondence optimization to two methods of producing corresponding PDMs: the PDMs computed via the SPHARM-PDM [10] method as well as PDMs optimized using the ShapeWorks software. To compare s-reps to the PDM data, we apply our measures of correspondence to the PDM implied by the s-rep spoke ends.

The rest of this paper lays out as follows. Sections 2 and 3 provide background information and a brief description of materials. Section 4 describes our method for interpolating continuous object interiors from discrete s-reps. Section 5 describes our method for improving the correspondence of s-rep models using an entropy-minimizing optimization. Section 6 gives results of applying the s-rep correspondence optimization, and section 7 discusses these results, gives conclusions, and describes future work.

2. Background

2.1. Entropy-based surface correspondence

Many shape representations require good correspondence over a population before being useful for statistical analysis. Manually constructing models with corresponding points is a difficult and time-consuming task. Instead some methods, such as SPHARM-PDM [10],

attempt to automatically generate representations with good correspondence based on geometry. Another approach is to take an existing model and attempt to optimize its correspondence by tightening a probability distribution. An entropy-based method for PDMs has been shown to improve correspondence across a population of objects [3]. This method takes an existing set of PDMs and shifts the points along the interpolated boundary to optimize an objective function which is a difference of two terms: a *geometry entropy* which measures how well some set of geometric features at a point match across the population, and a *regularity entropy* which measures how evenly the points are spread within each case. The idea is that the geometric probability distribution should be tightened while the probability of boundary point positions should be made more uniform on each case.

2.2. Skeletal Models

The s-rep [8] is a quasi-medial skeletal model that models not only an object's boundary but its interior as well. The s-rep is a collection of points sampled from a skeleton of an object which have associated vectors called *spokes* pointing from the skeleton to the object's boundary. This model is fit to an object via an optimization which requires the skeleton to be as close to medial as possible while remaining non-branching and requires the spokes to touch the object boundary and be nearly orthogonal to it [8]. As such, it captures not only object positions but also object widths and object boundary directions. These samples can then be interpolated to produce a continuous representation of an object's boundary and interior. This provides an object-relative coordinate system (u, v, τ) for the object's volume, where (u, v) parameterizes the skeleton and τ moves from the skeleton ($\tau = 0$) to the boundary ($\tau = 1$).

2.3. Composite Principal Nested Spheres

Because shapes in general, and s-reps in particular, have many features which do not naturally live in a flat Euclidean space, standard Euclidean statistical techniques such as principal component analysis (PCA) have proven inadequate for analyzing populations of s-reps. Instead, because many s-rep features live on spheres, such as the spoke directions (S^2) and the PDM formed from the n skeletal points (S^{3n}) [11], a method that can analyze data directly on spheres is preferred.

Analyzing spherical data can be done using a method called principal nested spheres (PNS) [12], which is analogous to PCA. As PCA repeatedly computes best-fitting subspaces of codimension 1 onto which the data is projected, PNS successively computes best-fitting subspheres, lowers the dimension by projection, and repeats. At each projection step, the signed geodesic residual of each data point to the subspace is recorded. It is reasonable to treat this collection of signed distances as a Euclidean data object and analyze it using standard methods. We call this process *Euclideanization*. By combining these Euclideanized features with the originally-Euclidean features and being careful to scale the Euclideanized features to make them commensurate, all of the data can now be analyzed using standard PCA. This method is called composite principal nested spheres (CPNS) [13].

3. Materials

The objects on which the methods described herein are evaluated consist of 31 lateral cerebral ventricles of neonates segmented from MRI. The subjects were at risk for schizophrenia or bipolar disorder.

Each ventricle was fit via the SPHARM-PDM [14] toolbox and 1002-point PDMs were extracted at corresponding latitudes and longitudes. An s-rep was fit [8] to the solid implied by each SPHARM.

4. Spoke Interpolation

4.1. Skeletal Mathematics

A continuous skeletal model describes an object interior by two functions: $\mathbf{p}(u_1, u_2)$, a 2D skeletal surface and $\mathbf{S}(u_1, u_2)$, a vector field pointing from the skeletal surface to the boundary. \mathbf{S} can be further decomposed into a product of two functions: $\mathbf{U}(u_1, u_2)$, a unit vector field pointing in the direction of \mathbf{S} and $r(u_1, u_2)$, a scalar distance function from the skeletal sheet to the object boundary.

An s-rep is a sampling of the continuous skeletal model. It consists of an $m \times n$ grid of samples from the skeletal surface with either two spokes (on the interior) or three (along the crest). The grid structure on the skeletal sheet forms a collection of quadrilaterals with a skeletal sample at each corner.

From this discrete representation we require a method to interpolate back to the continuous entity. Our method is a generalization to quasi-medial objects of the method for interpolating medial representations (m-reps) [15], which was based on the mathematics of medial structures [9].

We wish to interpolate a spoke $\mathbf{S}(u_1^*, u_2^*)$. If the sampled skeletal points have integer parameter values, (u_1^*, u_2^*) can be written as

$$(u_1^*, u_2^*) = (u_1^0 + \delta u_1, u_2^0 + \delta u_2); \quad \delta u_1, \delta u_2 \in [0, 1)$$

where (u_1^0, u_2^0) is the top left corner of the quad containing the desired value. From this, we obtain the equation

$$\mathbf{S}(u_1^*, u_2^*) = \mathbf{S}(u_1^0, u_2^0) + \oint_{(0,0)}^{(\delta u_1, \delta u_2)} \left(\frac{\partial \mathbf{S}}{\partial u_1} du_1 + \frac{\partial \mathbf{S}}{\partial u_2} du_2 \right) \quad (1)$$

for the spoke we wish to interpolate. Since $\mathbf{S}(u_1, u_2) = r(u_1, u_2)\mathbf{U}(u_1, u_2)$, the partial derivatives in equation 1 can be written as $\mathbf{S}_{u_i} = r\mathbf{U}_{u_i} + r_{u_i}\mathbf{U}$.

The radius derivatives r_{u_i} can be computed using the equation [16, 17]

$-r_{u_i} = \mathbf{p}_{u_i} \cdot \mathbf{U} + \mathbf{B}_{u_i} \cdot \mathbf{U}$, where \mathbf{B}_{u_i} are boundary derivatives, ($\mathbf{B} = \mathbf{p} + \mathbf{r}\mathbf{U} = \mathbf{p} + \mathbf{S}$ and $\mathbf{B}_{u_i} = \mathbf{p}_{u_i} + \mathbf{S}_{u_i}$). This equation is a correction to the compatibility condition for medial models [9] necessary because boundary and skeletal changes are not tied together as strongly in skeletal models as in medial ones. Using this expression for r_{u_i} , we obtain

$$\begin{aligned} \mathbf{S}_{u_i} &= r\mathbf{U}_{u_i} - (\mathbf{p}_{u_i} \cdot \mathbf{U} + (\mathbf{p}_{u_i} + \mathbf{S}_{u_i}) \cdot \mathbf{U})\mathbf{U} = r\mathbf{U}_{u_i} - (2\mathbf{p}_{u_i} + \mathbf{S}_{u_i})\mathbf{U}^T\mathbf{U} \\ &= (r\mathbf{U}_{u_i} - 2\mathbf{p}_{u_i}\mathbf{U}^T\mathbf{U})(\mathbf{I} + \mathbf{U}^T\mathbf{U})^{-1} \end{aligned} \quad (2)$$

This equation allows for the computation of the spoke derivative at any point (u_1^*, u_2^*) given the derivative of the skeletal surface \mathbf{p}_{u_i} and spoke direction \mathbf{U}_{u_i} at that point. Section 4.2 describes a method for computing \mathbf{p}_{u_i} , while 4.3 deals with computation of \mathbf{U}_{u_i} . Finally, 4.4 describes how a spoke is interpolated via numerical integration of equation 2.

4.2. Interpolation of the Skeletal Surface

Interpolation of the skeletal surface is done by fitting cubic Hermite patches to the quads of discrete samples which form the surface. This interpolation requires 16 values: the (4) corner points $\mathbf{p}(u_1^0, u_2^0)$, $\mathbf{p}(u_1^1, u_2^0)$, $\mathbf{p}(u_1^0, u_2^1)$, $\mathbf{p}(u_1^1, u_2^1)$, the (8) derivatives of each corner in both parameter directions, and the (4) second order mixed partial derivatives, which are set to 0. The control matrix H_c is thus

$$H_c = \begin{pmatrix} \mathbf{p}(u_1^0, u_2^0) & \mathbf{p}(u_1^1, u_2^0) & \mathbf{p}_{u_2}(u_1^0, u_2^0) & \mathbf{p}_{u_2}(u_1^1, u_2^0) \\ \mathbf{p}(u_1^0, u_2^1) & \mathbf{p}(u_1^1, u_2^1) & \mathbf{p}_{u_2}(u_1^0, u_2^1) & \mathbf{p}_{u_2}(u_1^1, u_2^1) \\ \mathbf{p}_{u_1}(u_1^0, u_2^0) & \mathbf{p}_{u_1}(u_1^1, u_2^0) & 0 & 0 \\ \mathbf{p}_{u_1}(u_1^0, u_2^1) & \mathbf{p}_{u_1}(u_1^1, u_2^1) & 0 & 0 \end{pmatrix}$$

Let $\mathbf{H}(s) = (H_1(s), H_2(s), H_3(s), H_4(s))$ and $\mathbf{H}'(s) = (H'_1(s), H'_2(s), H'_3(s), H'_4(s))$ where H_i are the cubic Hermite spline basis functions and H'_i are their derivatives. Computation at a point (u_1^*, u_2^*) inside a quad is given by

$$\mathbf{p}(u_1^*, u_2^*) = \mathbf{H}(\delta u_1) \cdot H_c \cdot \mathbf{H}(\delta u_2)^T.$$

Derivatives of the skeletal surface are computed by replacing the appropriate set of basis functions by their derivatives:

$$\mathbf{p}_{u_1}(u_1^*, u_2^*) = \mathbf{H}'(\delta u_1) \cdot H_c \cdot \mathbf{H}(\delta u_2)^\top; \mathbf{p}_{u_2}(u_1^*, u_2^*) = \mathbf{H}(\delta u_1) \cdot H_c \cdot \mathbf{H}'(\delta u_2)^\top \quad (3)$$

4.3. Estimation of Spoke Direction Derivatives via Quaternion Interpolation

In equation 2, the derivatives of the spoke direction vector field \mathbf{U}_{u_i} are needed. Because \mathbf{U} is a unit vector field, changes in \mathbf{U} are purely rotational. Thus, we choose a quaternion-based interpolation to compute these derivatives.

Each spoke direction at the corner of a quad is represented by a quaternion. A unit vector $\mathbf{U} = (U_x, U_y, U_z)$ is represented by the quaternion $\mathbf{q} = 0 + U_x i + U_y j + U_z k$. From the four spoke quaternions bounding a quad we can interpolate a unit quaternion in the quad interior by spherical linear interpolation (slerp) [18]. The quaternion that is $\lambda \in [0, 1]$ of the distance between \mathbf{q}_i and \mathbf{q}_{i+1} is given by $SL(\mathbf{q}_i, \mathbf{q}_{i+1}, \lambda) = \mathbf{q}_i (\mathbf{q}_i^* \mathbf{q}_{i+1})^\lambda$.

To achieve higher continuity across quad boundaries and thus a smoother surface, a higher order interpolation is desired. We use an extension of the cubic Bézier curve to the surface of a sphere called *squad* [19]. Analogously to the application of De Casteljau's algorithm to the computation of Bézier curves, *squad* can be computed in terms of slerp:

$$SQ(\mathbf{q}_i, \mathbf{q}_{i+1}, \mathbf{a}_i, \mathbf{a}_{i+1}, \lambda) = SL(SL(\mathbf{q}_i, \mathbf{q}_{i+1}, \lambda), SL(\mathbf{a}_i, \mathbf{a}_{i+1}, \lambda), 2\lambda(1-\lambda)) \quad (4)$$

where \mathbf{a}_i and \mathbf{a}_{i+1} are Bézier curve control points. Careful choice of these points ensures C^1 continuity across the \mathbf{q}_s [19][20]:

$$\mathbf{a}_i = \mathbf{q}_i \exp \left(-\frac{\log(\mathbf{q}_i^{-1} \mathbf{q}_{i+1}) + \log(\mathbf{q}_i^{-1} \mathbf{q}_{i-1})}{4} \right)$$

Differentiating the spoke interpolation formula SQ with respect to λ yields [20]

$$SQ'(\mathbf{q}_i, \mathbf{q}_{i+1}, \mathbf{a}_i, \mathbf{a}_{i+1}, \lambda) = SL(\mathbf{q}_i, \mathbf{q}_{i+1}, \lambda) \log(\mathbf{q}_i^* \mathbf{q}_{i+1}) \mathbf{g}_i(\lambda)^{2\lambda(1-\lambda)} + SL(\mathbf{q}_i, \mathbf{q}_{i+1}, \lambda) (\mathbf{g}_i'(\lambda))^{2\lambda(1-\lambda)} \quad (5)$$

where $\mathbf{g}_i(\lambda) = SL(\mathbf{q}_i, \mathbf{q}_{i+1}, \lambda) * SL(\mathbf{s}_i, \mathbf{s}_{i+1}, \lambda)$.

The derivative $\mathbf{U}_{u_1}(u_1^*, u_2^*)$ within a quad is computed by first using equation 4 to estimate

$\mathbf{U}(u_1^{-1}, u_2^{-2})$, $\mathbf{U}(u_1^0, u_2^{-2})$, $\mathbf{U}(u_1^1, u_2^{-2})$, and $\mathbf{U}(u_1^2, u_2^{-2})$ via the 4×4 surrounding grid spokes.

Equation 5 on the resulting quaternions then yields the desired derivative. The computation is similar for \mathbf{U}_{u_2} .

4.4. Spoke Computation via Integration of Derivatives

With the \mathbf{p}_{u_i} and \mathbf{U}_{u_i} values from sections 4.2 and 4.3, we can start from the quad corner (u_1^0, u_2^0) and integrate equation 2 using interval subdivision h to produce a spoke at $(u_1^h = u_1^0 + h\delta u_1, u_2^h = u_2^0 + h\delta u_2)$. Euler's method for the integration yields $\mathbf{S}(u_1^h, u_2^h) = \mathbf{S}(u_1^0, u_2^0) + h(\delta u_1 \mathbf{S}_{u_1}(u_1^0, u_2^0) + \delta u_2 \mathbf{S}_{u_2}(u_1^0, u_2^0))$. From (u_1^h, u_2^h) , we take another step towards (u_1^*, u_2^*) and iterate to (u_1^*, u_2^*) . Figure 2 shows results.

5. Correspondence

In this section we describe our method for improving s-rep correspondence through spoke shifting. Tightening the probability distribution of the geometric properties of s-reps in the population is the basic means of producing correspondence. If this is optimized alone, these features will tend to group together and produce highly irregular objects. This effect is avoided by also considering, for each training object, an entropy derived from various s-rep properties related to the regularity of the distribution of spokes throughout the s-rep. We improve s-rep correspondence by optimizing the weighted difference of these two entropy measures.

5.1. Spoke Shifting to Optimize Entropy

In order to optimize correspondence, spokes must be able to shift along the skeletal surface while respecting the boundary of the object. A spoke $\mathcal{S}(u_i, v_j)$, with $i, j \in \mathbf{Z}^+$, is shifted to a new position $\mathcal{S}(u_i + \delta u_i, v_j + \delta v_j)$ by interpolating the value of the spoke at $\mathcal{S}(u_i + \delta u_i, v_j + \delta v_j)$. By constraining $\delta u, \delta v \in (-0.5, 0.5)$, we ensure that the rectangular structure of the grid is kept intact. Spokes in the interior of the grid shift along the skeletal surface, while spokes along the fold of the object can only move on the crest curve.

The δu and δv by which the spokes are shifted are chosen to optimize an objective function which is the difference of two entropies: a geometry term E_{geo} which measures how well corresponding spokes match across a population of objects and a regularity term E_{reg} which measures how evenly spokes are distributed within each object. The following subsections discuss each of these terms in more detail.

5.2. Geometry Entropy

Correspondence in a population of s-reps is defined here as having a tight probability distribution of their geometric properties. These properties are the positions of the skeletal points, the spoke directions, and the spoke lengths. Because typical entropy measurements assume that the properties are Euclidean whereas s-rep features are not, we first Euclideanize and commensurate the object features as described in section 2.3 and detailed in [8]. After Euclideanization and commensuration, the probability distribution of these features can be estimated via standard Euclidean PCA, from which the entropy E_{geo} is computed.

5.3. Regularity Entropy

The regularity entropy is a measure of how uniformly distributed features are within each object. On an s-rep, using spoke interpolation the grid of skeletal positions can be connected to form curvilinear quadrilaterals with corresponding curvilinear quads on the object boundary. Figure 3 shows three examples of these quads with varying regularity.

S-rep regularity is measured on a number of features for each skeleton / boundary quad pair. We wish these features to be statistically independent while implying the geometrical properties of quad edge lengths, quad areas, and inter-quad-pair volumes, which are intuitively related to regularity. The features are as follows: horizontal and vertical quad edge lengths, average angle cosines of top-left and bottom-right corners of each interpolated sub-quad, and cosine of the swing of the quad normal between the two corners of the quad. The first two help to enforce regular shapes for the quads (tend towards rectangles) while the last penalizes highly curved surfaces (tends toward planar quads). These properties are computed separately for the skeleton and boundary quads. The similar features for each quad are combined (i.e., the boundary and skeletal horizontal edge lengths for that quad are concatenated) forming a tuple of that feature for that quad. These features are then combined from all quads into three groups: the horizontal edge lengths M_{hel} , the vertical edge lengths M_{vel} , and the angle cosines M_{cos} . If there are $2k$ copies of a feature per quad (k on the boundary and k on the skeleton) and there are n quads in the s-rep, this yields a $2k \times n$ feature vector.

From these three feature sets, the entropy measures E_{hel} , E_{vel} , and E_{cos} are derived (see section 5.4 for detail on the Entropy computation). The total measure of regularity for i^{th} object in the training set is $E_{reg}^i = E_{hel}^i + E_{vel}^i + E_{cos}^i$ due to the assumption of independence.

5.4. Optimization

As with the PDM correspondence method implemented in ShapeWorks, each of the entropies is computed using the assumption that the features follow a Gaussian distribution. The entropy of a d-dimensional Gaussian random variable X is $H(X) = \frac{1}{2} \left(d + d \ln(2\pi) + \sum_{i=1}^d \ln \lambda_i \right)$, where $\lambda_1, \lambda_2, \dots, \lambda_d$ are the non-zero eigenvalues of the covariance matrix of X . Because in many applications the covariance matrix will not have full rank, there will be a number of small eigenvalues that will disproportionately contribute to $H(X)$. We solve this by removing eigenvalues with contribution $\lambda_i / \sum \lambda_i$ lower than a threshold θ .

We optimize correspondence by minimizing

$$\arg \min_x \left(\omega E_{geo}(x) - \sum_{i=1}^N E_{reg}^i(x) \right) \quad (6)$$

where ω is a weight controlling the balance between probability distribution tightness and the regularity of each model and x is the shifting of each spoke in all of the models. The optimization is done via iterative, alternating applications of the NEWUOA [21] and one-

plus-one evolutionary [22] optimization algorithms. The choice of the parameters ω and θ was made via empirical evaluation.

6. Application & Results

Our method was first tested on a set of 80 synthetic lateral ventricle s-reps where all but one spoke were identical. After optimization, the distribution of the one moving spoke was greatly tightened.

In the following experiments we evaluate the method for improving s-rep correspondence on the set of 31 lateral ventricle shapes described in section 3. Unless otherwise noted, for these experiments $\omega = 4$ in equation 6 and $\theta = 0.01$. For each of the 31 lateral ventricles, in order to verify the accuracy of representation of the object after spoke shifting, we compute the Dice coefficient of the s-rep-implied boundaries before and after optimization. We obtain a mean Dice coefficient of 0.976, showing high agreement between the before and after surfaces.

6.1. Improvement of S-rep Correspondence

Figure 4 shows distributions of s-rep spokes before and after optimization. The s-reps after optimization have qualitatively better correspondence and maintain good regularity properties.

Figure 5 shows the cumulative variance explained by each eigenvalue from CPNS analysis on the s-rep population, i.e. of the s-rep features, before and after optimization. The two plots are similar, but after optimization the amount of variance captured in the first two eigenmodes has increased. The cumulative variance of the s-rep population is lower after optimization, with the sum of the eigenvalues being .0081 compared with .0085 before optimization.

Figure 6 shows the effect of the correspondence improvement on the s-rep population under three standard measures for evaluating correspondence [23, 24, 25]: compactness, generalization, and specificity. All of these measures are computed not on the s-rep directly, but on the PDM formed by taking two points for each spoke: the tail point on the skeletal sheet and the end point on the object boundary.

Compactness is a measure of tightness of a probability distribution, computed as the sum of the eigenvalues obtained via PCA. After correspondence optimization there is marked improvement in compactness.

Generalization is measured by computing a shape space on all but one of the training cases and computing the distance between the last shape and its projection onto this shape space. Generalization is computed on a set of PDMs in a leave-one-out fashion. PCA is applied to all but one case, forming a shape space. This final PDM is then projected onto the shape space given by the first M modes of variation. The distance between the PDM and its projection is computed as the average Euclidean distance between pairs of corresponding points. The optimized s-reps show slightly better generalization for low numbers of eigenvalues that would be used in a final model as indicated by the compactness result.

However, they show slightly worse generalization than before optimization at higher numbers of eigenvalues.

Specificity is a measure of average distance between random samples in the computed shape space with their nearest members of the data. Specificity is computed by using PCA on all PDMs to produce a probability distribution. Independent random samples are chosen from the distribution given by the first M modes. For each sample, the distance between it and the nearest original PDM is computed in the same way as for generalization. The optimized s-reps show noticeable improvement in specificity.

We also examined the effect of changing the weight ω on the results of the optimization. Figure 7 shows the correspondence measures on s-reps with different weights. In general, $\omega = 4$ is the best choice, showing the most improvement in compactness and specificity and the least worsening of generalization.

6.2. Comparison with PDM-based Methods

We compare the correspondence produced via optimized s-reps to the initial SPHARM-PDMs as well as PDMs optimized using the ShapeWorks software. For SPHARM-PDM and ShapeWorks, these measures are directly on the boundary PDMs. For the s-reps, use the PDMs mentioned previously. All PDMs have similar numbers of points, with the s-reps being interpolated to match the numbers of points in the SPHARM-PDMs. Figure 8 shows comparisons of the three methods using the three correspondence evaluation criteria. The s-rep-implied PDMs are scaled to match the scale of the others.

For compactness, the optimized s-reps show noticeably improved compactness over both of the PDM datasets. For generalization, the optimized s-reps are better than both for low numbers of eigenvalues but worse than the optimized PDMs at higher numbers. For specificity, the optimized s-reps again out-perform both PDM datasets.

7. Discussion

This paper presented novel methods for interpolating discrete s-reps into continuous objects and improving the correspondence of a population of s-reps. This interpolation is an extension of the interpolation of m-reps to the more general skeletal representation. The correspondence improvement is done via entropy minimization analogously to similar methods for PDMs. This optimization was shown to improve the correspondence of a population of lateral cerebral ventricle s-reps and s-rep-implied PDMs are shown to have correspondence comparable to or better than PDMs optimized via ShapeWorks.

The interpolation method presented here produces smoother and more accurate s-rep boundaries, as earlier methods for s-rep interpolation were based on the mathematics of medial structures and failed to correctly model the interaction between changes in the skeleton and the boundary when medial constraints are relaxed. Figure 9 shows a comparison of a boundary patch interpolated using both methods. The use of Euler's method for performing the integration steps will be replaced by a method with better speed and accuracy properties in future work.

The correspondence optimization results show good improvement in correspondence on the population of s-reps. Previous work has shown that s-reps without improved correspondence outperform PDMs in the statistical tasks of classification [7] and hypothesis testing [26] and we believe that using s-reps with improved correspondence will only further this advantage. In particular, the results on compactness show that the improved s-reps are a more efficient representation than the non-optimized s-reps used in previous studies, while the similar generalization and improved specificity suggest the improved s-reps could be more powerful for these tasks. Experiments to demonstrate this improvement are planned but are outside the scope of this paper.

The comparisons done to the PDM datasets are troublesome for several reasons. The use of PCA in computing probability distributions of shape representations, be they PDMs or s-reps, is not ideal because of the non-Euclidean nature of the data. Second, we compare purely boundary PDM-based methods (SPHARM-PDM and ShapeWorks) to PDMs implied by s-reps. While we use both boundary and skeletal points to leverage some of the interior correspondence s-reps provide, the use of point information without orientation or width information ignores many of the rich features that s-reps provide.

The current implementation of the correspondence optimization is in C++ and MATLAB takes 5 days to run on 31 cases on a 3.20GHz quad-core computer with 8GB of memory, with this time increasing with the number of cases being optimized. A reimplemention to speed up the optimization is ongoing.

Acknowledgements

The authors thank J.S. Marron and Sungkyu Jung for their consultation on statistical issues, Ilwoo Lyu for his help with the NEWUOA optimizer, and Dan Yang and Xiaohong Zhang for their support of Liyun Tu's work.

This work was supported in part by the National Institutes of Health (NIH) grants MH064065, MH070890, MH091351, MH091645, HD079124, and HD03110, and in part by the National Natural Science Foundation of China grants 91118005, 61173131 and 61402062.

References

- [1]. Van Kaick O, Zhang H, Hamarneh G, and Cohen-Or D, "A survey on shape correspondence," in *Computer Graphics Forum*, vol. 30, pp. 1681–1707, Wiley Online Library, 2011.
- [2]. Cootes TF, Taylor CJ, Cooper DH, and Graham J, "Active shape models-their training and application," *Computer vision and image understanding*, vol. 61, no. 1, pp. 38–59, 1995.
- [3]. Cates J, Fletcher PT, Styner M, Shenton M, and Whitaker R, "Shape modeling and analysis with entropy-based particle systems," in *Information Processing in Medical Imaging*, pp. 333–345, Springer, 2007.
- [4]. "Shapeworks." <http://www.sci.utah.edu/software/shapeworks.html>.
- [5]. Jung S and Qiao X, "A statistical approach to set classification by feature selection with applications to classification of histopathology images," *Biometrics*, vol. 70, no. 3, pp. 536–545, 2014. [PubMed: 24588775]
- [6]. Pizer SM, Hong J, Jung S, Marron JS, and Vicory J, "Relative statistical performance of s-reps with principal nested spheres vs. pdms," in *Shape 2014*, 2014.
- [7]. Hong J, Vicory J, Schulz J, Styner M, Marron JS, and Pizer SM, "Non-euclidean classification of medically imaged objects via s-reps," *Medical Image Analysis*, to appear.

- [8]. Pizer SM, Jung S, Goswami D, Vicory J, Zhao X, Chaudhuri R, Damon JN, Huckemann S, and Marron J, "Nested sphere statistics of skeletal models," in *Innovations for Shape Analysis*, pp. 93–115, Springer, 2013.
- [9]. Siddiqi K and Pizer SM, *Medial representations: mathematics, algorithms and applications*, vol. 37 Springer, 2008.
- [10]. Styner M, Oguz I, Xu S, Brechbu C, Pantazis hler, D., Levitt JJ, Shenton ME, and Gerig G, "Framework for the statistical shape analysis of brain structures using spharm-pdm," *The insight journal*, no. 1071, p. 242, 2006.
- [11]. Kendall DG, "Shape manifolds, procrustean metrics, and complex projective spaces," *Bulletin of the London Mathematical Society*, vol. 16, no. 2, pp. 81–121, 1984.
- [12]. Jung S, Dryden IL, and Marron J, "Analysis of principal nested spheres," *Biometrika*, vol. 99, no. 3, pp. 551–568, 2012. [PubMed: 23843669]
- [13]. Jung S, Liu X, Marron J, and Pizer SM, "Generalized pca via the backward stepwise approach in image analysis," in *Brain, Body and Machine*, pp. 111–123, Springer, 2010.
- [14]. "Spharm-pdm." <http://www.nitrc.org/projects/spharm-pdm>.
- [15]. Han Q, Pizer SM, and Damon JN, "Interpolation in discrete single figure medial objects," in *Computer Vision and Pattern Recognition Workshop, 2006. CVPRW'06. Conference on*, pp. 85–85, IEEE, 2006.
- [16]. Damon J, "Smoothness and geometry of boundaries associated to skeletal structures i: Sufficient conditions for smoothness," in *Annales de l'institut Fourier*, vol. 53, pp. 1941–1985, 2003.
- [17]. Damon JN personal correspondence, 2014.
- [18]. Shoemake K, "Animating rotation with quaternion curves," in *ACM SIGGRAPH computer graphics*, vol. 19, pp. 245–254, ACM, 1985.
- [19]. Shoemake K, "Quaternion calculus and fast animation, computer animation: 3-d motion specification and control," *Siggraph*, 1987.
- [20]. Dam EB, Koch M, and Lillholm M, *Quaternions, interpolation and animation*. Datalogisk Institut, Københavns Universitet, 1998.
- [21]. Powell MJ, "The newuoa software for unconstrained optimization without derivatives," in *Large-scale nonlinear optimization*, pp. 255–297, Springer, 2006.
- [22]. Styner M, Brechbuhler C, Szckely G, and Gerig G, "Parametric estimate of intensity inhomogeneities applied to mri," *Medical Imaging, IEEE Transactions on*, vol. 19, no. 3, pp. 153–165, 2000.
- [23]. Davies RH, *Learning shape: optimal models for analysing natural variability*. University of Manchester, 2002.
- [24]. Styner MA, Rajamani KT, Nolte L-P, Zsemlye G, Szekely G, Taylor CJ and Davies RH, "Evaluation of 3d correspondence methods for model building," in *Information Processing in Medical Imaging*, pp. 63–75, Springer, 2003.
- [25]. Davies RH, Twining CJ, Cootes TF, and Taylor CJ, "Building 3-d statistical shape models by direct optimization," *Medical Imaging, IEEE Transactions on*, vol. 29, no. 4, pp. 961–981, 2010.
- [26]. Schulz J, Pizer SM, Marron J, and Godtliebsen F, "Nonlinear hypothesis testing of geometrical object properties of shapes applied to hippocampi," 2013.

Highlights

- We develop a method for interpolating skeletal representations (s-reps).
- We optimize correspondence in an s-rep population via an entropy minimization.
- Interpolation is used to shift s-rep spokes to improve correspondence.
- We show improvement in s-reps under several correspondence measures.
- S-rep-implied PDMs have better correspondence than those from other methods.

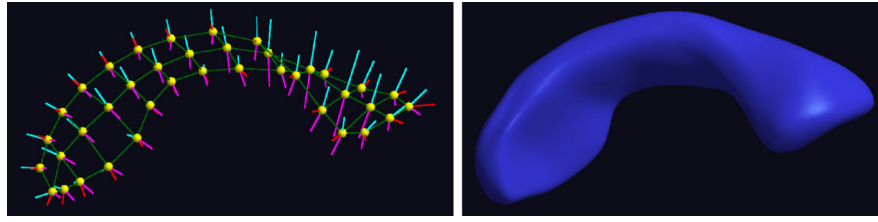


Figure 1:

An s-rep model (left) and its implied boundary (right). The yellow balls are samples of the skeletal surface. The cyan spokes point to the top of the object, the magenta to the bottom, and the red to the crest. The green lines show the grid structure of the skeletal samples.

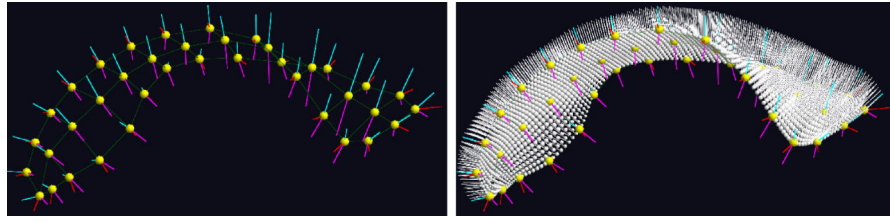


Figure 2:
A lateral ventricle s-rep and a dense interpolation of its top side spokes.

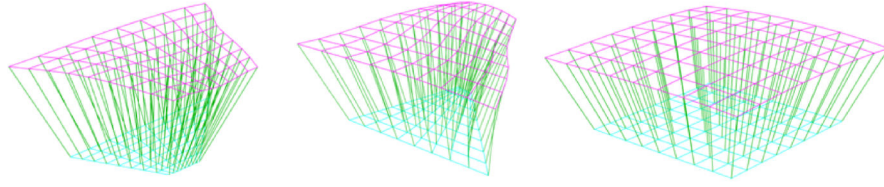


Figure 3: Three s-reps with varying regularity. The quads of the left and middle s-reps are non-regular while the right s-rep is highly regular.

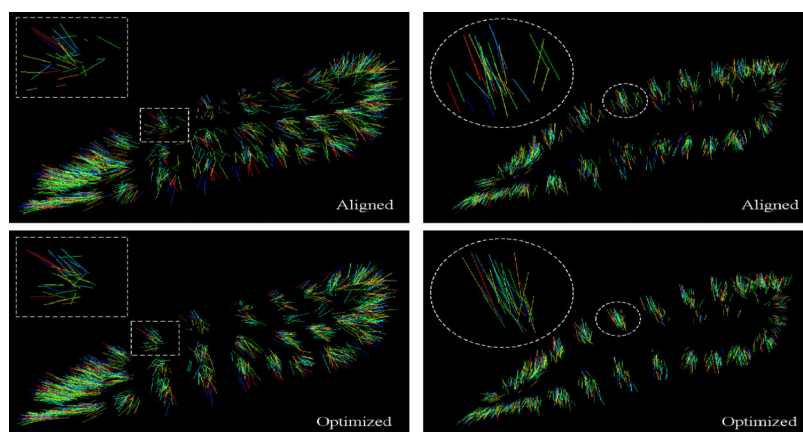


Figure 4: Distributions of spokes on the top (left) and crest side (right) of the s-reps before (first row) and after (second row) optimization. The groupings of spokes are noticeably tighter after optimization.

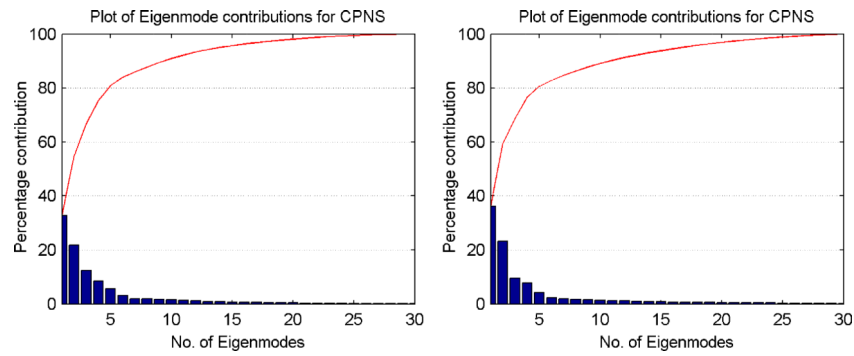


Figure 5: CPNS eigenvalues before (left) and after (right) correspondence optimization, as percent of total variance explained. The red curves show cumulative variance explained.

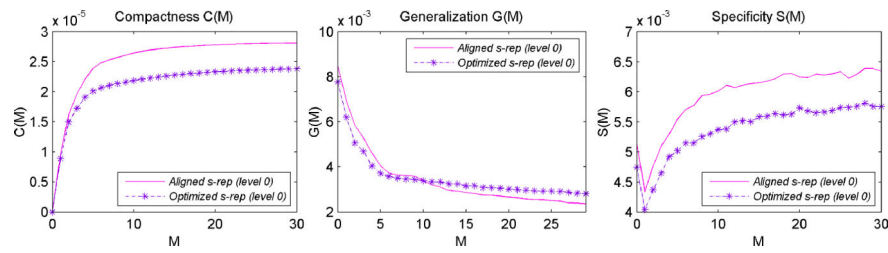


Figure 6: Plots of compactness (left), generalization (center), and specificity (right) on the original and optimized s-rep populations. These measures are functions of the number of eigenmodes (M) used to form the shape space.

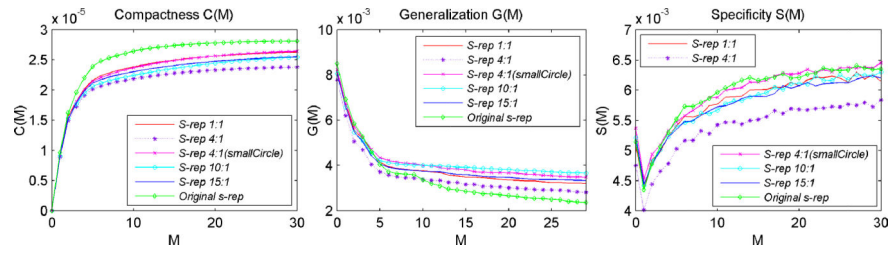
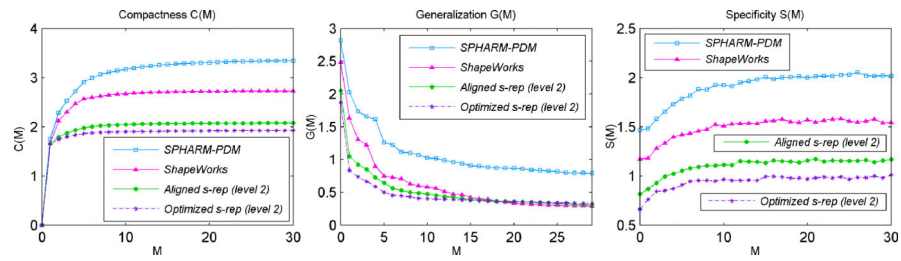


Figure 7: Effects of the weighting factor ω on the 31 lateral ventricle s-reps on compactness (left), generalization (middle), and specificity (right).

**Figure 8:**

Plots of compactness (left), generalization (middle), and specificity (right) measures of the original SPHARM-PDMs, the s-reps after correspondence optimization, and the PDMs optimized via ShapeWorks. These measures are functions of the number of eigenmodes (M) used to form the shape space.

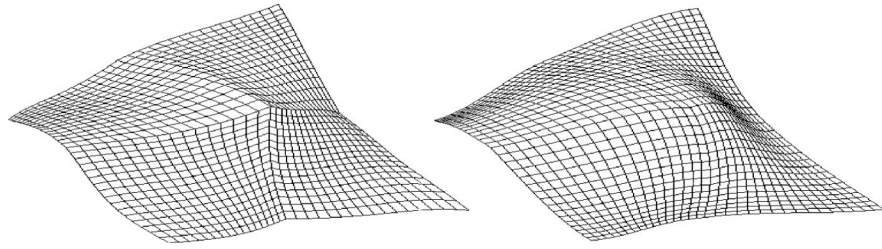


Figure 9:
Comparison between interpolations of an s-rep using medially-based (left) and skeletal-based (right) mathematics.