

Skeletal Shape Correspondence through Entropy

Liyun Tu, Martin Styner, Jared Vicory, Shireen Elhabian, Beatriz Paniagua, Juan C. Prieto, Dan Yang, Ross Whitaker and Stephen M. Pizer

Abstract—We present a novel approach for improving the shape statistics of medical image objects by generating correspondence of skeletal points. Each object’s interior is modeled by an *s-rep*, i.e., by a sampled, folded, 2-sided skeletal sheet with spoke vectors proceeding from the skeletal sheet to the boundary. The skeleton is divided into three parts: the up side, the down side and the fold curve. The spokes on each part are treated separately and, using spoke interpolation, are shifted along that skeleton in each training sample so as to tighten the probability distribution on those spokes’ geometric properties while sampling the object interior regularly. As with the surface/boundary-based correspondence method of Cates et al., entropy is used to measure both the probability distribution tightness and sampling regularity, here of the spokes’ geometric properties. Evaluation on synthetic and real world lateral ventricle and hippocampus datasets demonstrate improvement in the performance of statistics using the resulting probability distributions. This improvement is greater than that achieved by an entropy-based correspondence method on the boundary points.

Index Terms—Shape modeling and analysis, skeletal model, correspondence, lateral ventricle, hippocampus

I. INTRODUCTION

ESTABLISHING correspondence among similar shapes is essential for accurate statistical shape analysis [1, 2] in many biomedical applications such as understanding the anatomic structural differences in various stages of growth or disease [3-6]. This is often achieved by identifying a set of sparsely sampled and well-corresponding landmarks on organs or regions of interest across the shape instances.

Several representations of geometric models [6-11] have been proposed to describe anatomical structures. The most popular one is the surface/boundary-based point distribution

This work was supported in part by the National Institutes of Health (NIH) under Grants MH064065, MH070890, MH091351, MH091645, HD079124, and HD03110, and in part by the National Natural Science Foundation of China under Grants 91118005, 61173131 and 61402062, and in part by the National Institute of General Medical Sciences of the NIH under Grant P41GM103545. Asterisk indicates corresponding author.

*L. Tu is with the College of Computer Science, Chongqing University, Chongqing 400044, China and also with the Department of Computer Science, University of North Carolina at Chapel Hill, NC 27599 USA (e-mail: tuliyun@gmail.com).

M. Styner, J. Vicory and *S. M. Pizer are with the Department of Computer Science, University of North Carolina at Chapel Hill, NC 27599 USA. M. Styner also with the Departments of Psychiatry, University of North Carolina at Chapel Hill, NC 27599 USA (e-mail: {styner, jvicory, pizer}@cs.unc.edu).

S. Elhabian and R. Whitaker are with the Scientific Computing and Imaging Institute, and School of Computing, University of Utah, UT 84112 USA (e-mail: shireen@sci.utah.edu, whitaker@cs.utah.edu).

D. Yang is with the School of Software Engineering, Chongqing University, Chongqing 401331, China (e-mail: dyang@cqu.edu.cn).

B. Paniagua is with the Department of Psychiatry, University of North Carolina at Chapel Hill, NC 27599 USA (e-mail: bpaniagua@gmail.com).

J. C. Prieto is with the Center of Neurological Imaging, Brigham and Women’s hospital, MA 02115 USA (e-mail: juanprietob@gmail.com).

model (PDM) (e.g., [1, 12, 13]). Another well studied representation describes a shape using a set of parameterized basis functions such as spherical harmonics (SPHARM) [6, 8, 14], which defines an ambiguous correspondence and shown to be inadequate for some biomedical shapes that have non-uniform spherical parameterizations [15]. Deriving from the *m-reps* [16, 17], *skeletal* representation has proven powerful for shape analysis [11, 18-20] and biomedical applications [21-24] that need the rich features (e.g., local thickness and orientation) provided by the skeleton for quantifying the symmetric properties of organ shapes.

A key direction in the correspondence optimization research has been the computation of an objective function based on the determinant of the covariance matrix [25] and latter on the minimum description length (MDL) [26] to tighten the probability distribution of anatomically homologous points across a dataset (see Section II-B). In general, the MDL is equivalent to minimum entropy [21, 27]. According to this, an energy function that encodes the entropy of the geometric properties of the points in the ensemble of shapes (the *geometry entropy*) summed with the entropies of the point distribution on each shape’s boundary (the *regularity entropy*) was proposed in [1] to solve the point correspondence problem. This combined entropy allows them to optimize the geometric accuracy and the statistical simplicity of the shape model. However, in their method the shapes were described as surface/boundary-based PDMs, which ignore many of the higher order geometric features (e.g., local orientation) [18, 28].

In this paper, we adopt the skeletal model – termed as the *s-rep* [29] – which samples not just the object boundary but also its interior. The *s-rep* (Fig. 1) is formed by an approximately medial folded 2-sided skeletal sheet with a vector called a *spoke* proceeding from every skeletal point on the sheet to and approximately normal to the object’s boundary.




Fig. 1. Left: an example of *s-rep* for lateral ventricle that is sampled as a folded 3×13 skeletal sheet; right: object boundary (yellow) implied by that *s-rep*.

Our purpose is to build a novel correspondence method based on the aforementioned combined entropy to optimize the spoke geometry while keeping the spoke distribution in each object interior regular. Once correspondence is established, the *s-rep* provides an intrinsic coordinate system for the points in the object interior and near exterior that supports statistics on object geometry and object appearance.

The first contribution of this paper is a novel spoke sliding mechanism that shifts the discrete spokes to an interpolated spoke along the part of the skeletal sheet from which it proceeds. Spoke interpolation [24, 30, 31] is used to produce the new spoke at the shifted position.

The second contribution is a novel form of s-rep in which the skeletal positions of the up spokes, down spokes, and fold spokes are not necessarily shared (see Figure 4b). This new s-rep form is produced iteratively with the spoke shifting.

The third contribution is the spoke regularity properties that can effectively measure the regularity of the subregions bounded by the spokes representing an object.

The final contribution is to improve spoke correspondence by minimizing a weighted sum of the entropy of the spokes' geometric properties over the shape training population (the geometry entropy) and the within-object spokes' distributional entropy for each shape's interior (the regularity entropy).

Evaluated on synthetic and real word datasets consisting of lateral ventricles and hippocampi, the proposed method effectively improved the tightness and regularity of the spoke distribution and reduced shape variances. Also, when compared to entropy-based optimization of a boundary PDM, our method yields PDMs that are superior in specificity and compactness while maintaining generalization in the first few eigenmodes.

The remainder of this paper is organized as follows. Section II presents the shape representation schemes and overviews the previous related work along with common correspondence validation metrics which are adopted in this paper. Section III describes our correspondence method. Section IV details the experiments and presents measurements of the statistical performance of spoke optimization and boundary PDM optimization. Section V discusses this new correspondence method, including its possible improvements and applications.

II. BACKGROUND

A. Shape Representation Schemes

We describe three shape representations in use: SPHARM, PDM, and s-rep.

1) *SPHARM*: spherical harmonics coefficients as presented by Brechbühler et al. [8] form a global, fine scale boundary description that represents shapes of spherical topology. It describes a surface $\underline{x}(\theta, \varphi)$ using

$$\underline{x}(\theta, \varphi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l \underline{c}_l^m Y_l^m(\theta, \varphi), \quad (1)$$

where the basis functions $Y_l^m(\theta, \varphi)$, $-l \leq m \leq l$, of degree l and order m are defined on $\theta \in [0, \pi] \times \varphi \in [0, 2\pi)$ and where the three-dimensional coefficients \underline{c}_l^m are obtained by solving the least-squares problems in each spatial coordinate directions.

2) *PDM*: A PDM is a discrete set of points sampled manually or automatically from the object. Points with corresponding index in each object selected so as to be in correspondence across objects in the given ensemble. For statistical analysis, the points in each object are aligned via their moments to let all the examples lie in the same coordinate system.

Consider a set of N aligned object surface/boundary PDMs $\mathcal{S}_k = \{p_{ki}\}$, each with n points, where p_{ki} is a point on the k^{th} surface, $i \in (1, 2, \dots, n)$ and $k \in (1, 2, \dots, N)$. As shown by Kendall [32], scaling each center-of-mass-centered \mathcal{S}_k such

that $\sum_{i=1}^n |p_{ki}|^2 = 1$ can be seen as a projection onto the unit hypersphere \mathbb{S}^{3n-4} . Therefore, a boundary PDM can be understood as a concatenation of a log-transformed version of this scaling factor and this normalized tuple of points [33]; it abstractly lives on the manifold $\mathbb{R}^1 \times \mathbb{S}^{3n-4}$.

3) *S-rep*: In a discrete s-rep (Fig. 1), each up spoke shares a skeletal point with a down spoke, and fold spokes share a skeletal point with both an up and a down spoke. The object interior can be completely represented by interpolating the discrete s-rep into a continuous skeleton with a continuous field of spokes forming a continuous s-rep whose spokes fill the interior of the object. This representation does not rely on any inherent parameterization. Statistical analysis is applied to these sparse spokes [19, 24, 34].

For a new object, we obtain its s-rep by an optimization that fits the interpolated form of a previously defined s-rep template to the object. Because these fits involve common statistically generated modes of variation, followed by refinement, the fits already have a reasonable initial correspondence.

In this paper, we divide the skeleton into its three parts: up spokes proceeding from the skeletal sheet top, down spokes proceeding from the skeletal sheet bottom, and fold spokes. For the k spokes on each part, the set of spoke tails on the skeletal sheet forms a PDM to which the aforementioned PDM scaling strategy can be applied to yield a point on $\mathbb{R}^1 \times \mathbb{S}^{3k-4}$; the directional component of each spoke abstractly lives on the unit 2-sphere \mathbb{S}^2 . Therefore, the geometric properties of this s-rep abstractly lives on $\mathbb{R}^{k+1} \times \mathbb{S}^{3k-4} \times (\mathbb{S}^2)^k$. More information on this is provided in Supplementary Section I.

B. Previous Work on Correspondence

Cootes and Taylor [7] first posed the correspondence problem combined with Principal Component Analysis (PCA) on manually defined significance points in a 2D boundary PDM. Manually placing landmarks is error-prone, subjective and time-consuming, which makes this approach impractical for defining correspondence, the more so in 3D than 2D.

Various attempts have been made to solve the correspondence problem automatically [6, 12, 35, 36]. They can be categorized according to which of three alternative ways for building correspondence they follow. The first way is to manipulate surface correspondence simultaneously while computing geodesic deformation that is optimal according to an elastic Riemannian metric [13, 37-39]. The second way is to use a consistent, shape-sensitive method for fitting the object representation to the object, for each object in the training set [5, 29, 40-42]. The third way is to tighten the resulting probability distribution on the representation properties [1, 25-27, 43].

Srivastava et al. [13] and Jermyn et al. [37] introduced an elastic correspondence method that treats shape comparisons and surface correspondence in a joint manner. Kurtek et al. [10, 38] proposed a mechanism for computing geodesic paths between surfaces that are invariant to the parameterizations of surfaces and other shape-preserving transformations of surfaces, which was re-formulated and extended in [39] as a landmark-constrained correspondence method.

There are two extensively used methods based on fitting consistency. The first is applied to boundary PDMs based on SPHARM (see Section II-A1). The second is applied to s-reps and consequently to s-rep implied boundary PDMs.

The SPHARM-based method uses a PDM made from the vertices of a triangular tiling of the object boundary. It produces consistency in the areas of the tiles and in the orientations of the tiles relative to the parameterizing sphere. Kelemen et al. [40] used this method to describe a population of 3D hippocampal shapes. Inspired by their experiments, Gerig et al. [6] applied the same method to study the similarity of lateral ventricles. Styner et al. [41, 42] addressed the correspondence problem for anatomical objects through SPHARM-implied PDMs.

The s-rep-based method [5, 29] achieves correspondence by fitting a predefined template to each object instance. It achieves consistency either by deformation according to a common set of landmarks or by fitting each object over the coefficients of common eigenmodes produced by the PCA-like method called CPNS (Composite Principal Nested Spheres) [44, 45].

Correspondence based on tightening the probability distribution began with the work of Kotcheff and Taylor [25]. They posed the problem as one of finding the circle-based parameterization for each object in the training set that yields the tightest distribution. They measured tightness from the covariance matrix Σ on derived PDMs. Their measure was $\ln \det(\Sigma + \alpha I)$, where I is the identity matrix and α is a parameter that prevents the very thinnest modes (smallest eigenvalues) from disturbing the optimization process.

Davies et al. [26] proposed a different objective function based on Σ : MDL of the geometric representation [27]. In [12] they extended this approach to 3D statistical shape models based on a PDM description derived from a continuous parameterization. MDL uses an objective function $F = \sum_k \ln(\lambda_k + \alpha)$, where the λ_k are the first k dominant eigenvalues of Σ (principal variances) and α has the same effect as in Kotcheff's method. The regularity of the points on each object in this method was achieved via the deformation of the common parameterization.

The entropy $H(Z)$ of a d -dimensional Gaussian distribution on the multivariate Z is linear in the MDL objective function F with $\alpha = 0$ if all the principal variances are used. Thus for a fixed dimension, minimizing the MDL is equivalent to minimize the geometric entropy $H(Z)$ if all the principal variances are used and are offset by α . Realizing this, Cates et al. [1] recast the boundary PDMs correspondence problem in terms of entropy. This allowed them to express not only the geometric properties but also the regularity properties by entropy on each object. Their objective function was $Q = H_{Geometry}(Z) - \sum_{i=1}^N H_{Regularity}(Object_i)$. In both entropies, offsetting the principal variances by α was used. Q was minimized by shifting the points along each object boundary.

Other methods try to obtain correspondence through various geometric features. Brett et al. [46], Hill et al. [47] and Wang et al. [48, 49] proposed various shape properties (e.g., regions of high curvature) to establish point correspondences. Shi et al. [50] represented sulcal shapes as square-root velocity functions of continuous open curves and computed geodesics in the quotient space, trying to preserve important local geometry. Oguz [43] presented an extension of [1] by combining the features of point location with curvature and image based features. Her work illustrates how multiple modalities can be included in an entropy-based approach.

In this paper we present a correspondence method adopting

the aforementioned combined entropy for our s-rep to improve correspondence of the sparse spokes. That is, we describe the object as an s-rep and shift the spokes on the skeletal sheet to tighten the probability distribution of spoke geometry while keeping the spoke distribution in each object interior regular.

C. Correspondence Validation Metrics

A good correspondence method should generate a probability distribution on shape describable by few parameters; the probability distribution should be capable of representing objects that did not participate in training, and it should only generate instances that are similar to those in the training set. These three properties are respectively called compactness, generalization and specificity, all of which were first proposed by Davies' [51] and have been widely used in previous literature [12, 41, 43, 52] to evaluate the correspondence quality. For all three metrics, lower values are desirable.

The compactness metric of a model $C(M)$ evaluates the total variance of the model using $C(M) = \sum_{i=1}^M \lambda_i$, where λ_i is the i^{th} eigenvalue in the PCA model, M is the dimension of the object description. The standard error of $C(M)$ is $\sigma_{C(M)} = \sum_{i=1}^M \lambda_i \sqrt{2/N}$, where N is the number of objects in the training set. To compare shape descriptions with different number of features (e.g., PDMs with different number of points), we measure compactness in both variance and dimension using a per-dimension variant of the compactness metric as

$$C(M) = \frac{1}{K} \sum_{i=1}^M \lambda_i, \quad (2)$$

where K is the number of features representing each object.

The generalization ability of a model $G(M)$ measures its capability of representing unseen instances of the objects being studied. Consider an arbitrary object i ($i \in N$). A model is constructed by PCA using the remaining $N - 1$ objects (excluding the i^{th}). The model is used to reconstruct the i^{th} using M major modes of variation, and the reconstruction error is computed. $G(M)$ measures the average reconstruction error for all N objects as

$$G(M) = \frac{1}{N} \sum_{i=1}^N e(i, M), \quad e(i, M) = |X_i - X'_i(M)| \quad (3)$$

where X_i is the location array for the i^{th} object and $X'_i(M)$ is the location array for the reconstructed object for the i^{th} object using M modes. In practice, the standard error of $G(M)$ (denoted by $\sigma_{G(M)}$), $\sigma_{G(M)} = \sigma / \sqrt{N - 1}$, is usually computed to reason about the significance differences in generalization for different methods, where $\sigma = \sqrt{\frac{1}{N} \sum_{i=1}^N (e(i, M) - G(M))^2}$ is the sample standard deviation.

The specificity metric $S(M)$ of a model measures the degree to which the model generates objects that are similar to those in the training set. Generally, a PCA model is built from the training set (\mathcal{L}) and a population of m new instances (\mathcal{X}) are generated from this model using M modes. $S(M)$ measures the average distance error as

$$S(M) = \frac{1}{m} \sum_{i=1}^m d(i, M), \quad d(i, M) = |\Phi_i(M) - X_i| \quad (4)$$

where $\Phi_i(M)$ is the location array for the i^{th} instance in \mathfrak{L} . X_i is the location array for the object in \mathcal{L} that has the minimum distance to $\Phi_i(M)$. The standard error of $S(M)$ is given by $\sigma_{S(M)} = \sigma/\sqrt{m}$, where $\sigma = \sqrt{\frac{1}{m} \sum_{i=1}^m (d(i, M) - S(M))^2}$. A value of $m = 1000$ was used in this paper.

III. METHODS

Recall that our method optimizes the combined entropy in a training set of s-reps by sliding each of the sparse spokes (the thickened spokes of the same color in Fig. 2) on each object onto an interpolated spoke on the part of the skeletal sheet from which it proceeds. In this section, we elaborate on each of the method's components: a) spoke shifting through interpolation, b) regularity properties calculation contributing to regularity entropy, c) geometric properties calculation contributing to geometry entropy, d) objective function and optimization.



Fig. 2. Discrete s-reps of two lateral cerebral ventricles, each represented as a 3×13 sparse grid of spokes (the colored thickened lines). The thin gray spokes are the interpolated spokes describing the object interior. The color indicates the initial correspondence before any shifting using entropies.

A. Spoke Shifting through Interpolation

As previously discussed and as illustrated in Fig. 3a, an s-rep is made up of three regions of spokes: up, down and fold. The spoke shifting within each region is done independently. For an $m \times n$ s-rep, there are $(m - 2) \times (n - 2)$ up (or down) spokes emanating from an interior grid position; the remainder emanate from an exterior grid position (the fold). The interior up and down spokes should be allowed to shift on the interior of the respective region via the spoke interpolation in [31]. The spokes emanating from the fold should be allowed to shift only on the fold via the spoke interpolation method in [24]. The spokes emanating from the corner of the grid are not allowed to shift. Fig. 3b visualizes the shifting of the up region.

The spoke interpolation is based on an object-related intrinsic coordinate system, in which each *quad* (bounded by four neighboring spokes) on the skeletal sheet is projected as a unit grid. The sparse spokes are given successive integer grid coordinates (u, v) , with u parameterizing the s-rep grid's long axis and v parameterizing its short axis (Fig. 4a).

To optimize the aforementioned combined regularity and geometry statistics, each sparse spoke $S(u, v)$ is slid to the shifted spoke at the interpolated position $(u + \Delta u, v + \Delta v)$, where $\Delta u, \Delta v \in (-1, 1)$. Neighboring spokes are constrained so that their tails cannot cross on the skeletal sheet.

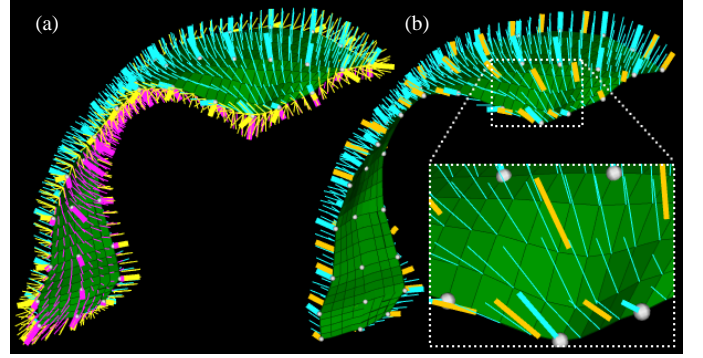


Fig. 3. (a) Each object is modeled with three regions: an up (cyan), a down (magenta) and a fold region (yellow). The green grids represent the skeletal sheet. (b) The up skeletal sheet with the original sparse spokes shown as thickened cyan lines; the interpolated spokes shown as thin cyan lines; the shifted spokes shown in thickened yellow. The thickened cyan spokes $S(u, v)$ are shifted to the thickened yellow position $S(u + \Delta u, v + \Delta v)$ by a small step $(\Delta u, \Delta v)$, and interpolating to that position.

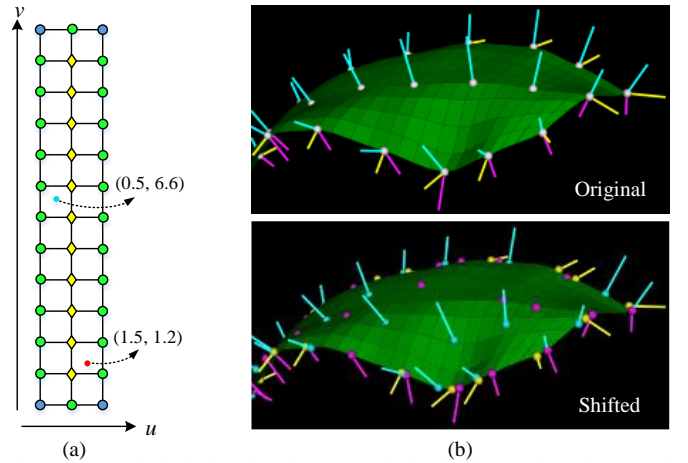


Fig. 4. (a) A skeletal sheet has its 3×13 skeletal points projected as a unit grid. The diamonds and balls respectively denote the interior and the exterior grid positions. Each interpolated spoke has coordinate $(u + \Delta u, v + \Delta v)$ (e.g., the spoke at the red dot is $(1.5, 1.2)$, and at the cyan dot $(0.5, 6.6)$). (b) Part of the s-rep before and after shifting. The shared skeletal points (white balls in original s-rep) split and shift to a new place with its spoke. The cyan ball is the skeletal pt. of the up spoke, magenta (down) and yellow (fold).

In general, each movable spoke has restriction on where it can shift to, according to the following criteria:

- 1) For the interior up and down spokes (diamonds in Fig. 4a), $\Delta u \in (-1, 1)$ and $\Delta v \in (-1, 1)$.
- 2) The spokes on the exterior of the grid (green balls in Fig. 4a), be they up, down or fold spokes, shift to the same types of spokes in an adjacent quad edge by changing only Δu or Δv . That is, for these spokes, the shifting is $\Delta u = 0, \Delta v \in (-1, 1)$ or $\Delta u \in (-1, 1), \Delta v = 0$.
- 3) For the corner spokes (blue balls in Fig. 4a), the shifting is $\Delta u = 0$ and $\Delta v = 0$.

After applying the above shifting independently to each region, the resulting s-rep may well not have any two spokes sharing a skeletal position (see Fig. 4b).

To produce correspondence, the spoke shifting is optimized iteratively until convergence is achieved. Whereas the interpolation used in the shifting is always based on the original s-rep; the regularity and geometric properties to be discussed in Sections III-B,C are always based on the shifted s-rep.

B. Regularity Properties

To measure the regularity of the spoke distribution of each s-rep in the training set, the grid of sparse spokes is formed into quads on both the skeletal sheet (connecting four neighbor spoke tails, shown as green balls and black grids in Fig. 5a-c) and the object boundary (connecting four neighbor spoke tips, shown as green or red sheets in Fig. 5a). At the fold there are curvilinear segments on the skeleton rather than quads. The regularity of the spokes is defined based on a tuple of features for each skeletal quad and its corresponding boundary quad. For each such quad-pair certain features can be understood to be approximately statistically independent and yet imply quad side lengths, areas, and inter-quad-pair volume. These are horizontal edge lengths, vertical edge lengths, average angle cosines over the sub-quads, and cosine of the normal swing between the upper-left and bottom-right corners. The length and angle properties are computed separately for the skeleton and boundary surfaces. Similar properties are computed for the region bounded by the fold curve on the skeleton and the corresponding fold region on the boundary (Fig. 5e).

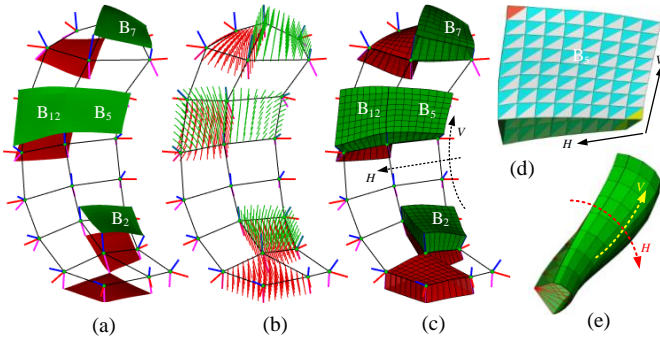


Fig. 5. Visualization of an 3×8 s-rep. (a) This s-rep has 24 up/down spokes (blue/magenta lines), 18 fold spokes (red lines). These spokes form 14 skeletal quads for the up/down region, each with a corresponding boundary quad (e.g., B_5 denotes a boundary quad). (b-c) Each quad is divided into $(2^3)^2$ sub-quads. (d) Each sub-quad in B_5 is further subdivided into two triangles. (e) An example of the curvilinear segments on the skeleton and its corresponding fold region on the boundary. H and V denote the horizontal and vertical directions respectively. More subdivision is visualized in Supplementary Fig. 1 for up/down region and Supplementary Fig. 2 for fold region.

All the quad edges are curved. As visualized in Fig. 5, we measure the lengths and angles by subdividing the quads and summing the lengths and averaging the angles measured from each subdivision. Moreover, quad subdivision into triangles handles the fact that the subdivision is curved. The normal swing is an approximation counting from the upper-left corner sub-triangle (e.g., the orange triangle in Fig. 5d) to the bottom-right corner sub-triangle (e.g., the yellow triangle).

For each training s-rep, all the resulting measurements are analyzed to produce three regularity entropy terms (E_{cos} , E_{hel} and E_{vel}) from the following three approximately statistically independent feature sets: the angle cosines, the horizontal edge lengths, and the vertical edge lengths.

C. Geometric Properties

Tightening the probability distribution on the geometric properties of the s-reps in the training set is the basic means of producing correspondence. These properties include the

positions of the skeletal points, the spoke radii, and the spoke directions. As with the regularity properties calculation, the geometric properties for the up, down and fold regions are also calculated separately and then combined across all regions. The geometric properties of the spokes in each region are computed in the same way.

A complication is that the entropy formulas used to measure tightness assume the properties are Euclidean, i.e., the Pythagorean theorem applies, but unfortunately this does not hold for many of the aforementioned geometric properties. We deal with this by Euclideanizing the relevant properties via Principal Nested Spheres (PNS) analysis [44] and making them commensurate before applying PCA and thence the entropy formulas. The commensurated form of each scale factor of the k^{th} s-rep, denoted by γ^k , is $\bar{\gamma} \times \log(\gamma^k/\bar{\gamma})$, where $\bar{\gamma}$ is the geometric mean of γ^k and $k = 1, 2, \dots, N$. Each Euclideanized feature resulting from the PNS on the scaled skeletal points is commensurated by multiplying it with $\bar{\gamma}$. The commensurated, Euclideanized form of the spoke radii r_i^k , are similarly $\bar{r}_i \times \log(r_i^k/\bar{r}_i)$, where \bar{r}_i is the geometric mean of the radii for the i^{th} spoke, $i = 1, 2, \dots, n$. PNS is applied to each spoke direction, which abstractly lives on S^2 , and the Euclideanized features are commensurated by multiplication with \bar{r}_i .

The result of this process is a collection of commensurate zero-mean features that can be concatenated to form a Euclidean feature tuple describing each s-rep. The sample covariance matrix derived from these tuples is used to form the geometry entropy E_{geo} , as will be described in Section III-D.

D. Energy Function and Optimization

Euclideanization and commensuration of the geometric features of the i^{th} s-rep, $i = 1, 2, \dots, n$, yields a k -dimensional feature tuple X_i that approximately follows a multivariate Gaussian distribution, $X \sim N(\mu, \Sigma)$, where μ is the mean tuple, Σ is a $k \times k$ covariance matrix and n is the number of s-reps in the training set [44]. This formulation can also be applied to the regularity properties after they have been made zero-mean and commensurate. The entropy of such an X is

$$H(x) = \frac{d}{2} + \frac{d}{2} \ln 2\pi + \frac{1}{2} \sum_{i=1}^d \ln \lambda_i \quad (5)$$

where d is the number of non-zero eigenvalues of Σ and $\lambda_1, \lambda_2, \dots, \lambda_d$ are the non-zero eigenvalues of Σ .

In our application $d < k$, with the effect that computed eigenvalues beyond the d^{th} are small and contain almost pure noise. Including these in the sum contributes dominating negative terms of large magnitude to $H(x)$. Recall that Cates et al. method used a constant factor to prevent the smallest modes (those with smallest eigenvalues), which mainly constitute effects of noise, from disturbing the optimization process. In our method, we handle this problem by removing those eigenvalues with contribution $(\lambda_i/\sum \lambda_i)$ smaller than a prior threshold ϑ and correspondingly lowering d . This approach is applied separately to compute the geometry entropy E_{geo} and each of the three regularity entropies (E_{cos} , E_{hel} and E_{vel}) for each object. For all the regularity entropies $\vartheta = 1\%$ is used. Because the dimensionality of the geometric properties is high, for the geometry entropy $\vartheta = 0.1\%$ is used.

The regularity entropy E_{reg}^i of the i^{th} s-rep is computed as

$$E_{reg}^i = E_{cos}^i + E_{hel}^i + E_{vel}^i \quad (6)$$

where E_{cos} indicates the angle cosine entropy. The overall regularity entropy of the N training s-reps is $E_{reg} = \sum_{i=1}^N E_{reg}^i$.

The correspondence is established by solving the optimization problem given by

$$\underset{x}{\operatorname{argmin}} f(x) = \{x \mid f(x) = \omega E_{geo} - E_{reg}\} \quad (7)$$

where ω is the weight used to balance the tightness and regularization and x is the collection of $(\Delta u, \Delta v)$ that shift the spokes. For the results reported later, we use the NEWUOA optimizer [53] to minimize (7) and use the one-plus-one evolutionary optimizer [54] to avoid local optima. However, other optimizers are possible.

IV. EVALUATION AND RESULTS

A series of experiments were performed to evaluate the proposed correspondence method. We first introduced the data sets in use. Second, we set the parameters involved in our objective function. Then we applied the proposed method to different data sets to investigate the tightness and regularity of the spoke distribution and the statistical parameters that were extensively used to measure the correspondence quality (see Section II-C). Finally, we compared our method with two of the state-of-the-art PDM-based methods, namely, SPHARM-PDM (spherical harmonics point distribution model) [15] and ShapeWorks [55] which is an up-to-date, enhanced open-source distribution of Cates et al. method [1]. Our program was implemented in C++ and Matlab. All experiments presented here were done on 64-bit 8GB 4-core PC.

A. Data Sets

Three types of data sets were employed for the results reported in this paper: synthetic objects and real world lateral ventricles and hippocampi.

1) *Synthetic Objects*: A set of 80 synthetic lateral ventricles s-reps with all spokes identical except one that was shifted by a small random distance were produced to evaluate how the correspondence optimization affects the spoke variance.

2) *Lateral Ventricles*: We were provided set of lateral ventricles that were semi-automatically segmented from magnetic resonance imaging (MRI) images in neonates [56]. Each ventricle was preprocessed using SPHARM-PDM and deformed to an initial skeletal model (manually defined by an expert) using thin plate spline registration. We selected 31 resulting ventricle s-reps for our tests presented here.

3) *Hippocampi*: We were provided s-reps fitted to a set of binary images of the hippocampi that were segmented from the MRI scans from normal control datasets in a schizophrenia study [5]. We randomly selected 40 of them for our tests.

For the latter two data sets, a Procrustes alignment based on the boundary points and their corresponding skeletal points implied by s-rep spokes was performed before applying the proposed correspondence method.

B. Parameter Selection

We first determined the best interpolation level (∂) for computing the regularity and geometric properties. Then we

studied how to determine the weight to best balances the tightness and regularization entropies.

1) Interpolation Level

The regularity properties described in Section III-B were computed based on subdivision of the curved quad edges into 2^∂ linear pieces. We sought the ∂ that balances the accuracy and efficiency by investigating the changes of the skeletal and boundary quad areas and of the 3D region volume bounded by the corresponding quad-quad (or line-quad for fold region) pair and the average time used for interpolating to each region of 10 different s-reps with $\partial = 0, 1, \dots, 8$ on the real lateral ventricles.

As shown in Supplementary Fig. 3, both areas and the volume change significantly as we increase ∂ up to 2, but there is little change when we increase ∂ from 2 to any high level. That is, $\partial = 0$ or $\partial = 1$ is inappropriate for the regularity properties computation; while $\partial = 2$ or 3 can achieve adequate accuracy. Since $\partial = 3$ takes three times the time of $\partial = 2$ (Supplementary Table I), we use $\partial = 2$: each edge is divided into 4 sub-edges, each quad is divided into 16 sub-quads.

2) The Weight for the Tightness and Regularization

Consider the weight ω on E_{geo} in (7). For different datasets, different values of ω yield best correspondence quality, as measured through s-rep implied boundary points. As an example, Supplementary Fig. 4 compares the statistical performance via specificity, generalization, and compactness measures as described in Section II-C for the real lateral ventricles using different ω . S-rep spoke tips ($\partial = 0$) were collected and used as boundary PDMs for these measurements. The results show that $\omega = 4$ obtains the best performance.

We suggest that $\omega = E_{reg}/E_{geo}$ is a reasonable initial value to try for any particular dataset, where E_{reg} and E_{geo} are calculated from the input, aligned s-reps; then one would switch ω up and down to investigate if better results can be obtained. This approach yielded weights $\omega = 30$, $\omega = 4$ and $\omega = 13$, respectively used for our synthetic objects, real lateral ventricles and hippocampi.

C. Comparisons between Aligned and Optimized S-reps

We applied the proposed method to the three aforementioned datasets and compared the entropies, shape variances and correspondence quality measurements between the aligned, non-optimized s-reps (referred to as *aligned s-reps*) and the correspondence-optimized s-reps from our method (referred to as *optimized s-reps*). S-rep spoke tips at subdivision level 0 ($\partial = 0$) were used as the implied boundary points (PDMs) for the statistical measurements in this subsection.

1) Synthetic Objects

We began with a predefined template s-rep from a lateral ventricle (Fig. 6a). Each synthetic s-rep was formed from the template by moving only one particular fold spoke. The spoke was moved to an interpolated fold position by a small distance according to a Gaussian distribution (Fig. 6b-c). As depicted in Fig. 6c-d, after our program is applied the spokes agreement improved in location as well as lengths. Similar results were observed in Supplementary Fig. 5.

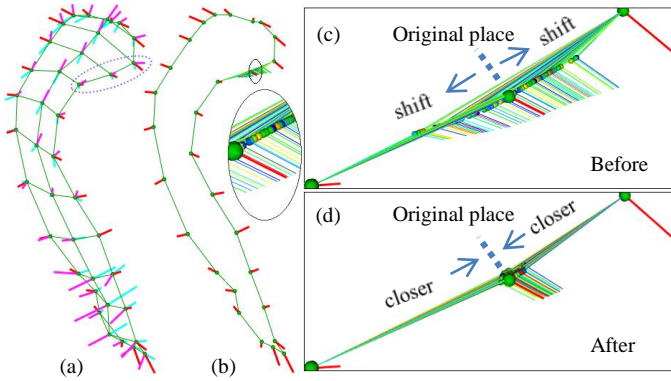


Fig. 6. (a) The template s-rep; (b) the fold curve (green curve) and fold spoke (red line) of the template and the 80 synthetic s-reps on top of each other, with the corresponding fold spoke dispersed from the original place (the red line pointing out from the green ball in oval) by a step with $\sigma = 0.4$; (c-d) the corresponding spaced spokes distribution before and after correspondence optimization, respectively. The color denotes each case in the training s-reps.




Fig. 7. Left: entropies and the objective function ($f(x)$ in (7)) during the iterations for the 80 synthetic objects; right: the changing of E_{geo} of the up region for the real hippocampi.

The evolution of entropy values during the optimization are illustrated in Fig. 7 (left panel). During the iterations E_{geo} decreases from -10.47 to -38.64. E_{reg} changes little (from the beginning of 531.76 to the end of 533.95); nevertheless, the 3D region volumes bounded by the fold curve on skeleton and the corresponding quad on boundary associated with this fold spoke have their variances in the ratio $\frac{\sigma_{optimized}}{\sigma_{original}} = 0.16$; and as shown in Table I, the standard deviation of each regularity property decreases. All three metrics of the shape model performance improved (Supplementary Fig. 6) in a similar way as that graphed in Fig. 8.

Table I. Standard deviation of the quad regularity properties associate with that fold spoke across the training set before and after optimization. Vertical edge length on skeleton (σ_{vel_s}), vertical edge length on boundary (σ_{vel_b}), angle cosine of the upper-left corner ($\sigma_{cos_{ul}}$), angle cosine of the bottom-right corner ($\sigma_{cos_{br}}$) and angle cosine of the normal swing ($\sigma_{cos_{ns}}$).

	σ_{vel_s}	σ_{vel_b}	$\sigma_{cos_{ul}}$	$\sigma_{cos_{br}}$	$\sigma_{cos_{ns}}$
Original	0.0059	0.0059	0.0417	0.0469	0.1414
Optimized	0.0007	0.0008	0.0208	0.0323	0.0262

2) Hippocampi

For the hippocampi dataset the entropies of up, down and fold regions were minimized separately. In each region, E_{reg} increased and E_{geo} decreased (e.g., Fig. 7, right panel for up region) in a similar pattern (Supplementary Table II).

The correspondence quality between the optimized s-reps and the aligned s-reps are compared in Fig. 8. The error bars

graphed on each line show the significance differences in each metric. The specificity (Fig. 8a) and generalization (Fig. 8b) measures for the s-rep-implied PDMs are based on the s-rep shape space. The compactness measures were computed both from the s-reps native geometric properties (Fig. 8c) and then from the s-rep-implied PDMs (Fig. 8d), whether on geometric properties or PDMs, were all computed via CPNS. As detailed in Section II-C, for all three measures lower values mean better performance.

The optimized s-reps are superior in compactness and specificity. In generalization, the aligned s-reps are superior for a number of eigenmodes ($M > 10$). However, the compactness graphs suggest that fewer than 10 eigenmodes would be enough for statistical applications, and in this range the generalization of the optimized s-reps are slightly superior.

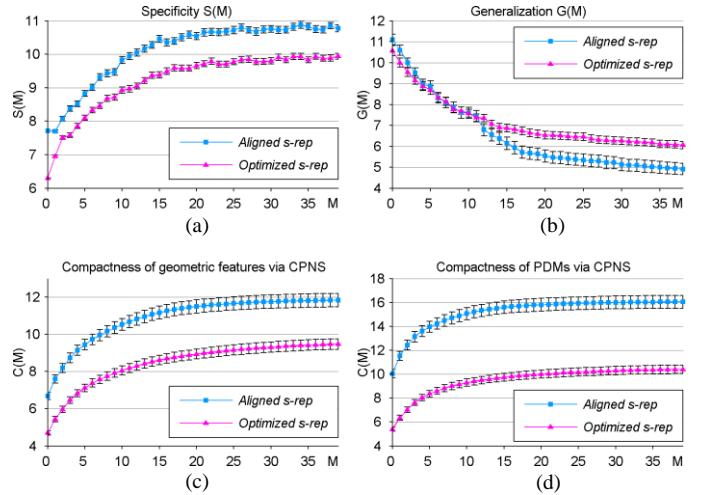


Fig. 8. Comparisons between optimized and aligned s-reps for the hippocampi. (a) specificity; (b) generalization; (c-d) compactness with s-reps geometric properties and with s-rep implied PDMs, respectively. For all three metrics, lower values are desirable.

3) Lateral Ventricles

For the lateral ventricle dataset the entropies of all three regions are shown in Table II. In all regions, E_{reg} increases (reflecting improved regularity) while E_{geo} decreases (reflecting a tightened probability distribution). The entropies of the up and the down regions have similar values because of the symmetry of the brain lateral ventricle, but E_{geo} of the up region decreases more than that of the down region, possibly because the surface of the up region has lower curvature and its spoke geometries are less tight at the beginning.

Table II. Regularity entropy (E_{reg}) and geometry entropy (E_{geo}) of each region before (R1) and after (R2) optimization.

	Up region		Down region		Fold region	
	E_{reg}	E_{geo}	E_{reg}	E_{geo}	E_{reg}	E_{geo}
R1	290.0	-35.7	305.4	-36.6	214.4	-35.4
R2	334.8	-41.2	348.2	-38.0	226.3	-50.2

As shown in Table III, optimization yields smaller geometric eigenvalues in all regions. The dominant portion of this decrease is in the first six eigenvalues (Supplementary Fig. 7). The portion of total variance captured by the two dominant

eigenvalues increased (Table IV), suggesting that the optimized s-reps capture more variance in fewer principal directions.

Table III. Total variances ($\sum \lambda_i$) of the lateral ventricle s-reps

	Up region	Down region	Fold region
Aligned	0.0034	0.0031	0.0019
Optimized	0.0029	0.0028	0.0015

Table IV. Contribution of the first two dominant s-rep eigenvalues before and after optimization.

	Up region (%)		Down region (%)		Fold region (%)	
	Before	After	Before	After	Before	After
λ_1	36.73	39.10	38.04	42.04	31.32	37.17
λ_2	22.67	25.08	21.39	23.00	23.65	27.06

The statistical shape model performance of optimized s-reps compared with aligned s-reps is laid out in Fig. 9. As with Fig. 8 in Section IV-C2, $S(M)$ (Fig. 9a) and $G(M)$ (Fig. 9b) measures for the s-reps are based on s-rep shape spaces. $C(M)$ was computed via CPNS from both s-reps features (Fig. 9c) and then from s-rep-implied PDM features (Fig. 9d).

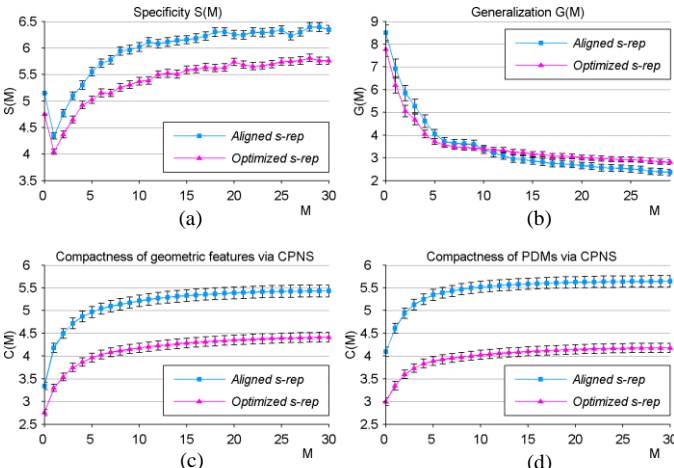


Fig. 9. Comparisons between optimized and aligned s-reps for lateral ventricles. (a) specificity; (b) generalization; (c-d) compactness with s-reps geometric features and with s-rep-implied PDMs ($\delta = 0$), respectively.

As with the previous hippocampi dataset, the optimized s-reps are also superior to the aligned s-reps in compactness and specificity. In generalization, the optimized s-reps are superior in the first 9 eigenmodes, which would be enough for statistical applications according to the compactness graph.




Fig. 10. The hippocampi (first row) and lateral ventricles (second row) model mean and ± 3 standard deviations in two eigenmodes. The shapes are generated from boundary PDMs implied by the optimized s-reps level 2 ($\delta = 2$) spoke tips.

D. Comparisons between S-reps and PDM-based Methods

As a further validation of the proposed method, this experiment compares the statistical performance among the optimized s-rep implied PDMs, the SPHARM-PDM and the ShapeWorks on real lateral ventricles and hippocampi datasets.

On inspection, all the resulting s-reps appear to be of good quality and imply high quality boundary PDMs. As an example, Fig. 10 shows the surface mesh of the mean PDM implied by the optimized s-reps as well as its deformations along the first two dominant s-rep implied PDM eigenmodes. Each main eigenmode describes some plausible pattern of shape changes observed in the population.

The PDM-based evaluations for the s-reps involved in our comparison were based on spoke boundary points only (named B_PDM) and on boundary and skeletal points (named BS_PDM), for each optimized s-rep at interpolation level 2. This yields a B_PDM of 1218 points and a BS_PDM of 2212 points for a 3×13 lateral ventricle s-rep; and of 738 and 1332 points, respectively for a 3×8 hippocampus s-rep. The SPHARM-PDMs used subdivision level 10 and SPHARM degree 11, which yields a PDM with 1002 points. The ShapeWorks used these SPHARM-PDMs as the input and optimized them with scaling off. All PDMs were scaled to lie in a same space.

The comparisons of statistical measurements among these PDMs are laid out in Fig. 11 (top row for lateral ventricles and bottom row for hippocampi).

For both data sets, the specificity and compactness measures on our optimized s-rep implied PDMs are noticeably superior to both the SPHARM-PDM and the ShapeWorks methods. The generalization measure on s-reps is better than both methods for low numbers of eigenmodes (M), e.g., $M < 15$ for lateral ventricles and $M < 5$ for hippocampi, but worse than the optimized PDMs at higher numbers especially for hippocampi. This is probably because the hippocampi have small curvature and most of the shape variance in the training population located on the boundary so that pure boundary PDM optimization performs better.

The BS_PDM gains in all measures over the B_PDM for both objects, suggesting that the skeletal points capture some interior correspondence. This observation is consistent with [22, 29] that the use of point information alone ignores many of the higher order geometric features that s-reps provide, such as orientation and width.

As shown in Supplementary Fig. 8, PCA-based compactness measures on the PDMs, albeit not properly applicable to the aforementioned non-Euclidean shape representations (see

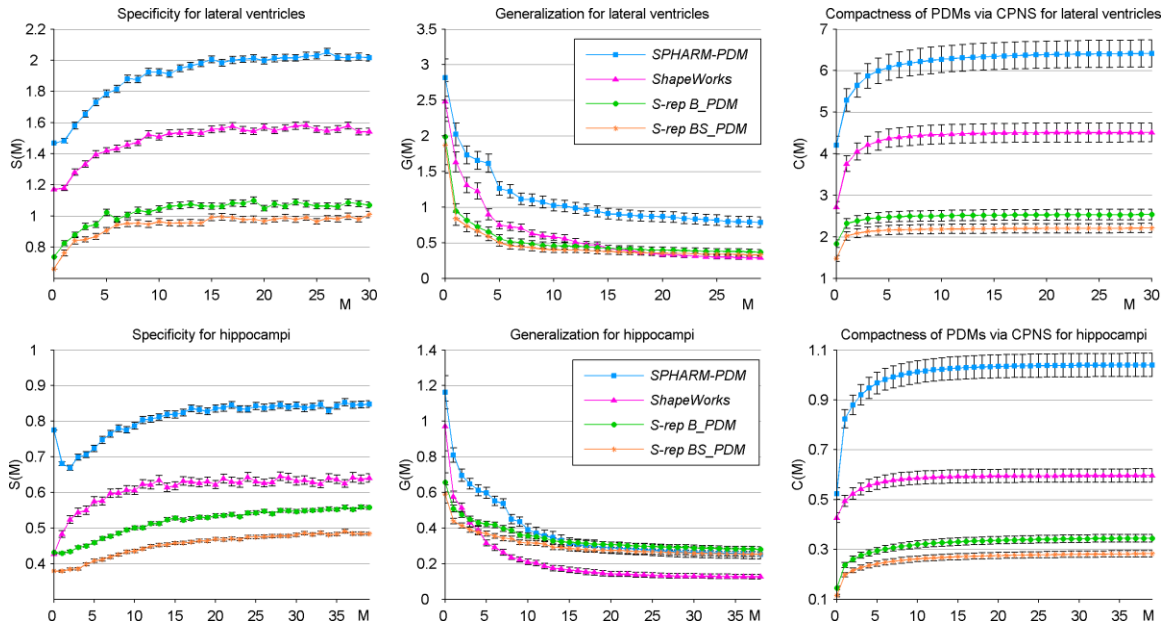


Fig. 11. Comparisons among the optimized s-rep implied PDMs, the SPHARM-PDM and the ShapeWorks on lateral ventricles (top row) and hippocampi (bottom row). The compactness for all models were computed via CPNS. All figures share the same legend as located in the middle column.

Section III-C and Supplementary Section I), led to the same conclusion as those based on CPNS.

V. CONCLUSION AND FUTURE WORK

In this work we proposed a novel group-wise optimization of skeletal properties to establish an enhanced s-rep correspondence. The proposed method represents each object in the training set as an s-rep, whose spokes are shifted along its skeletal part using spoke interpolation in each sample so as to tighten the probability distribution on those spokes' geometric properties while sampling the object interior regularly. The correspondence is established by minimizing an objective function that balances entropy derived from geometric properties and entropy derived from regularity properties. All these properties are computed from s-rep spokes.

This method effectively lowers the entropies and improves the correspondence of the spokes. It noticeably tightens the distribution of the corresponding spokes over the training set and disperses the spokes of each s-rep regularly. And it noticeably reduces the total variance of the spoke geometries. Moreover, the proposed method yields models with improved model properties as measured via generalization, specificity and compactness. Surprisingly, on our two objects sets, the result of our entropy-based correspondence optimization on the skeleton showed superior boundary point statistics even when compared to an entropy-based boundary correspondence methodology.

In this work, we evaluated the established s-rep correspondence with respect to model shape probability distributions. Main applications of shape models though include classification, hypothesis testing, segmentation via shape statistics, and registration in statistically generated shape spaces. These applications were not evaluated here and future research is needed to show the improvement of our s-rep

correspondence method for these purposes.

ACKNOWLEDGEMENT

The authors thank Sungkyu Jung and J. S. Marron for developing and providing the code for the CPNS method. The authors also thank Ilwoo Lyu for sharing information about the NEWUOA optimizer and Jörn Schulz for providing the hippocampi s-reps.

REFERENCES

- [1] J. Cates, P. T. Fletcher, M. Styner, M. Shenton, and R. Whitaker, "Shape modeling and analysis with entropy-based particle systems," in *Proc. Inf. Proc. Med. Imag.*, 2007, vol. 20, pp. 333-345.
- [2] M. Styner, and G. Gerig, "Three-dimensional medial shape representation incorporating object variability," in *IEEE Conf. Comput. Vis. Pattern Recognit.*, Kauai, HI, 2001, pp. 651-656.
- [3] L. Shen, H. A. Firpi, A. J. Saykin, and J. D. West, "Parametric surface modeling and registration for comparison of manual and automated segmentation of the hippocampus," *Hippocampus*, vol. 19, no. 6, pp. 588-595, 2009.
- [4] T. Heimann, and H.-P. Meinzer, "Statistical shape models for 3D medical image segmentation: A review," *Med. Image Anal.*, vol. 13, no. 4, pp. 543-563, 2009.
- [5] J. Schulz, S. M. Pizer, J. S. Marron, and F. Godtlielsen, "Nonlinear hypothesis testing of geometric object properties of shapes applied to hippocampi," *J. Math. Imaging Vis.*, 2015. to be published.
- [6] G. Gerig, M. Styner, M. E. Shenton, and J. A. Lieberman, "Shape versus size: improved understanding of the morphology of brain structures," in *Proc. Med. Image Comp. Comput.-Assist. Intervent.*, Netherlands, 2001, vol. 2208, pp. 24-32.
- [7] T. F. Cootes, and C. J. Taylor, "Combining point distribution models with shape models based on finite element analysis," *Image Vision Comput.*, vol. 13, no. 5, pp. 403-409, 1995.
- [8] C. Brechbühler, G. Gerig, and O. Kübler, "Parametrization of closed surfaces for 3-D shape description," *Comput. Vis. Image Understand.*, vol. 61, no. 2, pp. 154-171, 1995.
- [9] P. A. Yushkevich, H. Zhang, and J. C. Gee, "Continuous medial representation for anatomical structures," *IEEE Trans. Med. Imag.*, vol. 25, no. 12, pp. 1547-1564, 2006.

- [10] S. Kurtek, E. Klassen, J. C. Gore, Z. Ding, and A. Srivastava, "Elastic geodesic paths in shape space of parameterized surfaces," *IEEE Trans. Patt. Anal. Mach. Intell.*, pp. 1717-1730, 2012.
- [11] S. Bouixa, J. C. Pruessner, D. L. Collins, and K. Siddiqi, "Hippocampal shape analysis using medial surfaces," *NeuroImage*, vol. 25, no. 4, pp. 1077-1089, 2005.
- [12] R. H. Davies, C. J. Twining, T. F. Cootes, and C. J. Taylor, "Building 3-D statistical shape models by direct optimization," *IEEE Trans. Med. Imag.*, vol. 29, no. 4, pp. 961-981, 2010.
- [13] A. Srivastava, E. Klassen, S. H. Joshi, and I. Jermyn, "Shape analysis of elastic curves in euclidean spaces," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 33, no. 7, pp. 1415-1428, 2011.
- [14] D. Meier, and E. Fisher, "Parameter space warping: shape-based correspondence between morphologically different objects," *IEEE Trans. Med. Imag.*, vol. 21, no. 1, pp. 31-47, 2002.
- [15] M. Styner, I. Oguz, S. Xu, C. Brechbühler, D. Pantazis, J. J. Levitt, M. E. Shenton, and G. Gerig, "Framework for the statistical shape analysis of brain structures using SPHARM-PDM," *Open Science Workshop at MICCAI - Insight J.*, no. 1071, pp. 242-250, 2006.
- [16] S. M. Pizer, D. S. Fritsch, P. A. Yushkevich, V. E. Johnson, and E. L. Chaney, "Segmentation, registration, and measurement of shape variation via image object shape," *IEEE Trans. Med. Imag.*, vol. 18, no. 10, pp. 851-865, 1999.
- [17] K. Siddiqi, and S. M. Pizer, *Medial representations: mathematics, algorithms and applications*, Springer, 2008.
- [18] P. A. Yushkevich, and H. G. Zhang, "Deformable modeling using a 3D boundary representation with quadratic constraints on the branching structure of the Blum skeleton," in *Proc. Inf. Proc. Med. Imag.*, Berlin, 2013, pp. 280-291.
- [19] L. Tu, M. Styner, J. Vicory, B. Paniagua, J. C. Prieto, D. Yang, and S. M. Pizer, "Skeletal shape correspondence via entropy minimization," in *Proc. SPIE Med. Imag.*, Florida, 2015, vol. 9413.
- [20] J. Martínez, N. Pla, and M. Vigo, "Skeletal representations of orthogonal shapes," *Graphical models*, vol. 75, pp. 189-207, 2013.
- [21] M. Styner, J. A. Lieberman, D. Pantazis, and G. Gerig, "Boundary and medial shape analysis of the hippocampus in schizophrenia," *Med. Image Anal.*, vol. 8, pp. 197-203, 2004.
- [22] J. Hong, J. Vicory, J. Schulz, M. Styner, J. S. Marron, and S. M. Pizer, "Classification of medically imaged objects via s-rep statistics," unpublished.
- [23] P. A. Yushkevich, "Continuous medial representation of brain structures using the biharmonic PDE," *NeuroImage*, vol. 45, pp. S99-S110, 2009.
- [24] J. C. Prieto, "Multiparametric organ modeling for shape statistics and simulation procedures," Ph.D. dissertation, Inst. Natl des Sci. App. de Lyon, Univ. de Lyon, France, 2014.
- [25] A. C. W. Kotcheff, and C. J. Taylor, "Automatic construction of eigenshape models by direct optimization," *Med. Image Anal.*, vol. 2, no. 4, pp. 303-314, 1998.
- [26] R. H. Davies, C. J. Twining, T. F. Cootes, J. C. Waterton, and C. J. Taylor, "A minimum description length approach to statistical shape modeling," *IEEE Trans. Med. Imag.*, vol. 21, no. 5, pp. 525-537, 2002.
- [27] H. H. Thodberg, "Minimum description length shape and appearance models," in *Proc. Inf. Proc. Med. Imag.*, 2003, pp. 51-62.
- [28] S. M. Pizer, J. Hong, S. Jung, J. S. Marron, J. Schulz, and J. Vicory, "Relative statistical performance of s-reps with principal nested spheres vs. PDMs," in *Proc. Shape 2014 - Symp. of Stat. Shape Models and Appl.*, Delémont, 2014, vol. 10, pp. 11-13.
- [29] S. M. Pizer, S. Jung, D. Goswami, J. Vicory, X. Zhao, R. Chaudhuri, J. N. Damon, S. Huckemann, and J. S. Marron, "Nested sphere statistics of skeletal models," *Innovations for Shape Analysis: Models and Algorithms, Part I, LNCS*, Berlin: Springer, 2013, pp. 93-115.
- [30] J. Damon, "Smoothness and geometry of boundaries associated to skeletal structures I: Sufficient conditions for smoothness," *Annales de l'Institut Fourier*, vol. 53, no. 6, pp. 1941-1985, 2003.
- [31] L. Tu, J. Vicory, S. Elhabian, B. Paniagua, J. C. Prieto, R. Whitaker, M. Styner, and S. M. Pizer, "Entropy-based correspondence improvement of interpolated skeletal models," unpublished.
- [32] D. G. Kendall, "Shape manifolds, procrustean metrics, and complex projective spaces," *Bulletin of the London Mathematical Society*, vol. 16, no. 2, pp. 81-121, 1984.
- [33] S. Jung, X. Liu, J. S. Marron, and S. M. Pizer, "Generalized PCA via the backward stepwise approach in image analysis," in *Proc. Int. Symp. on the 25th Anniv. of McGill Univ. Centre for Intell. Mach., J. Angeles et al. (Eds.) Adv. in Intell. and Soft Comput. in Brain, Body and Mach.*, 2010, vol. 83, pp. 111-123.
- [34] J. Schulz, S. Jung, S. Huckemann, M. Pierrynowski, J. S. Marron, and S. M. Pizer, "Analysis of rotational deformations from directional data," *J. Comput. Graph. Stat.*, vol. 24, no. 2, pp. 539-560, 2015.
- [35] T. F. Cootes, and C. J. Taylor, "Statistical models of appearance for medical image analysis and computer vision," in *Proc. SPIE Med. Imag.*, 2001, vol. 4322, pp. 236-248.
- [36] D. Shen, E. H. Herskovits, and C. Davatzikos, "An adaptive-focus statistical shape model for segmentation and shape modeling of 3-D brain structures," *IEEE Trans. Med. Imag.*, vol. 20, no. 4, pp. 257-270, 2001.
- [37] I. Jermyn, S. Kurtek, E. Klassen, and A. Srivastava, "Elastic shape matching of parameterized surfaces using square root normal fields," in *Proc. Eur. Conf. Comput. Vis.*, Italy, 2012, pp. 804-817.
- [38] S. Kurtek, E. Klassen, Z. Ding, M. J. Avison, and A. Srivastava, "Parameterization-Invariant shape statistics and probabilistic classification of anatomical surfaces," in *Proc. Inf. Proc. Med. Imag.*, Germany, 2011, pp. 147-158.
- [39] S. Kurtek, A. Srivastava, E. Klassen, and H. Laga, "Landmark-guided elastic shape analysis of spherically-parameterized surfaces," *Comput. Graph. Forum*, vol. 32, no. 2, pp. 429-438, 2013.
- [40] A. Kelemen, G. Szekely, and G. Gerig, "Elastic model-based segmentation of 3d neuroradiological data sets," *IEEE Trans. Med. Imag.*, vol. 18, no. 10, pp. 828-839, 1999.
- [41] M. A. Styner, K. T. Rajamani, L.-P. Nolte, G. Zsemlye, G. a. Székely, C. J. Taylor, and R. H. Davies, "Evaluation of 3D correspondence methods for model building," in *Proc. Inf. Proc. Med. Imag.*, 2003, vol. 18, pp. 63-75.
- [42] M. Styner, G. Gerig, J. Lieberman, D. Jones, and D. Weinberger, "Statistical shape analysis of neuroanatomical structures based on medial models," *Med. Image Anal.*, vol. 7, no. 3, pp. 207-220, 2003.
- [43] I. Oguz, "Groupwise shape correspondence with local features," Ph.D. dissertation, Dept. of Comput. Sci., Univ. of North Carolina at Chapel Hill, Chapel Hill, 2009.
- [44] S. Jung, I. L. Dryden, and J. S. Marron, "Analysis of principal nested spheres," *Biometrika*, vol. 99, no. 3, pp. 551-568, 2012.
- [45] J. Damon, and J. S. Marron, "Backwards principal component analysis and principal nested relations," *J. Math. Imaging Vis.*, vol. 50, no. 2, pp. 107-114, 2014.
- [46] A. D. Brett, and C. J. Taylor, "A method of automated landmark generation for automated 3D PDM construction," *Image Vision Comput.*, vol. 18, no. 9, pp. 739-748, 2000.
- [47] A. Hill, C. J. Taylor, and A. D. Brett, "A framework for automatic landmark identification using a new method of nonrigid correspondence," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 22, no. 3, pp. 241-251, 2000.
- [48] Y. Wang, B. S. Peterson, and L. H. Staib, "Shape-based 3D surface correspondence using geodesics and local geometry," in *IEEE Conf. Comput. Vis. Pattern Recognit.*, 2000.
- [49] Y. Wang, J. Shi, X. Yin, X. Gu, T. F. Chan, S.-T. Yau, A. W. Toga, and P. M. Thompson, "Brain surface conformal parameterization with the Ricci flow," *IEEE Trans. Med. Imag.*, vol. 31, no. 2, pp. 251-264, 2012.
- [50] S. H. Joshi, R. P. Cabeen, A. A. Joshi, B. Sun, I. Dinov, K. L. Narr, A. W. Toga, and R. P. Woods, "Diffeomorphic sulcal shape analysis on the cortex," *IEEE Trans. Med. Imag.*, vol. 31, no. 6, pp. 1195-1212, 2012.
- [51] R. H. Davies, "Learning shape: optimal models for analysing natural variability," Ph.D. dissertation, Dept. of Imag. Sci. and Biomed. Eng., Univ. of Manchester, England, 2002.
- [52] R. Xu, X. Zhou, Y. Hirano, R. Tachibana, T. Hara, S. Kido, and H. Fujita, "Particle system based adaptive sampling on spherical parameter space to improve the MDL method for construction of statistical shape models," *Comput. Math. Method. M.*, pp. 1-9, 2013.
- [53] M. J. D. Powell, "The NEWUOA software for unconstrained optimization without derivatives," *Large-Scale Nonlinear Optim.*, vol. 83, pp. 255-297, 2006.
- [54] M. Styner, C. Brechbühler, G. Szekely, and G. Gerig, "Parametric estimate of intensity inhomogeneities applied to MRI," *IEEE Trans. Med. Imag.*, vol. 19, no. 3, pp. 153-165, 2000.
- [55] R. Whitaker. (2014, Sep). The ShapeWorks open-source software. [Online]. Available: <http://www.sci.utah.edu/software/shapeworks.html>.
- [56] B. Paniagua, A. Lyall, J.-B. Berger, C. Vachet, R. M. Hamer, S. Woolson, W. Lin, J. Gilmore, and M. Styner, "Lateral ventricle morphology analysis via mean latitude axis," in *Proc. SPIE Med. Imag.*, 2013, vol. 8672, pp. 20-27.