Statistics of Pose and Shape in Multi-Object Complexes using Principal Geodesic Analysis

Martin Styner[1,2], Kevin Gorczowski[1], Tom Fletcher[1],
Ja Yeon Jeong[1], Stephen M. Pizer[1], Guido Gerig[1,2]

The University of North Carolina
Departments of [1]Computer Science and [2]Psychiatry
Chapel Hill, NC 27599

Abstract. A main focus of statistical shape analysis is the description of variability of a population of geometric objects. In this paper, we present work in progress towards modeling the shape and pose variability of sets of multiple objects. Principal geodesic analysis (PGA) is the extension of the standard technique of principal component analysis (PCA) into the nonlinear Riemannian symmetric space of pose and our medial m-rep shape description, a space in which use of PCA would be incorrect. In this paper, we discuss the decoupling of pose and shape in multi-object sets using different normalization settings. Further, we introduce new methods of describing the statistics of object pose using a novel extension of PGA, which previously has been used for global shape statistics. These new pose statistics are then combined with shape statistics to form a more complete description of multi-object complexes. We demonstrate our methods in an application to a longitudinal pediatric autism study with object sets of 10 subcortical structures in a population of 20 subjects. The results show that global scale accounts for most of the major mode of variation across time. Furthermore, the PGA components and the corresponding distribution of different subject groups vary significantly depending on the choice of normalization, which illustrates the importance of global and local pose alignment in multi-object shape analysis.

1   Introduction

Statistical shape modeling and analysis [1, 2] is emerging as an important tool for understanding anatomical structures from medical images. Statistical shape modeling's goal is to construct a compact and stable description of the mean and variability of a population. Principal Component Analysis (PCA) is probably the most widely used procedure for generating shape models of variability. These models can provide understanding for processes of growth and disease observed in neuroimaging [3].

Clinical applications favor a statistical shape modeling of multi-object sets rather than one of single structures outside of their multi-object context. Neuroimaging studies of mental illness and neurolocal disease, for example, are interested in describing group differences and changes due to neurodevelopment or neurodegeneration. These processes most likely affect multiple structures rather than a single one. A description of the change of the object set might help to explain underlying neurobiological processes affecting brain circuits. Whereas Tsai et al. [4] and Yang et al. [5] describe statistical

object modeling by level-sets, we propose explicit deformable shape modeling with a sampled medial mesh representation called m-rep, introduced by Pizer et al [6].

Deformable shape models represent the underlying geometry of the anatomy and then use a statistical analysis to describe the variability of that geometry. Several different geometric representations other than m-reps have been used to model anatomy, such as landmarks [7], dense collection of boundary points [8], or spherical harmonic decompositions [9]. Another shape variability approach focuses on the analysis of deformation maps [10, 3, 11, 12]. A fundamental difficulty in statistical shape modeling is the high dimensionality of the set of features with a relatively small sample size, typically in the range of 20 to 50 in neuroimaging studies. This problem is even more evident for modeling sets of multiple objects. The analysis of transformation fields has to cope with the high dimensionality of the transformation, which renders the computation of the PCA basically unstable in respect to the training set. Adding or removing a single subject from the training set results in strikingly different principal components.

In most shape modeling approaches, the underlying geometry is parameterized as an Euclidean feature space. As Davis and Joshi [13] note, the space of diffeomorphism is a curvilinear space. Treating it as an Euclidean space is a linear simplification of a higher-dimensional, curvilinear problem. For medial descriptions, as well as for descriptions of pose, the feature space clearly contains elements of a non-Euclidean vector space. These features need to be parameterized in a nonlinear Riemannian symmetric space. We will use curved statistics for these parameters with modeling major modes of deformations via principle geodesic analysis (PGA) [14], a nonlinear extension of PCA.

This paper summarizes work in progress towards an efficient and compact representation and statistical analysis of sets of objects. We choose the sampled medial m-rep representation and a statistical framework based on Riemannian metrics. The driving application is a longitudinal pediatric neuroimaging study.

2   Methodology

This research is driven by the challenge to describe the shape statistics of a set of 3-D objects. Whereas analysis of single shapes is well advanced and has been described extensively using a variety of shape parametrization techniques, the research community does not yet have access to tools for statistical modeling and analysis of sets of objects.

Estimating Variability of Multi-object Sets: In linear space, variability of parameterized objects can be described by principle component analysis (PCA) of spherical harmonics [9] or point distribution models (PDM) [8]. However, a linear PCA cannot describe object rotations and the modeling cannot be extended to model points and normals. Non-linear modeling is achieved by principle geodesic analysis (PGA), developed by Fletcher et al. [14]. PGA extends linear PCA into nonlinear space using "curved statistics" and is a natural generalization of PCA for data that are parameterized as curved manifolds. To recall, the intrinsic mean of a collection of points $x_1, \cdots, x_N$ on a Riemannian manifold M is the Fréchet mean $\mu = argmin \sum_{i=1}^{N} d(x, x_i)^2$, where $d(.,.)$ denotes Riemannian distance on M. Whereas PCA in $\mathbb{R}^3$ generates linear subspaces that maximize the variance of projected data, geodesic manifolds are images of linear subspaces under the exponential map and are defined as the manifolds that

maximize projected variance. Principle geodesics can be found by a recursive gradient descent or with an approximation by the log map and a linear PCA in the tangent space of the map (please see [15] for details). An important fact is that PGA can be used with parametrization schemes that include point locations, scale, and angle parameters.
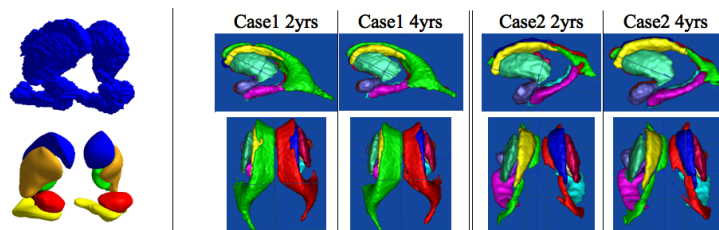
Fig. 1. Visualization of the 10 selected deep brain structures. Left: Binary voxel objects (top) and implied solid surface of medial atoms (bottom). Right (with lateral ventricles): Left side view (top row), Top view (bottom row). The inter-subject shape differences clearly are larger than the longitudinal differences, which seem quite small.

Object Representation by a Mesh of Medial Samples: Medial representations represent an alternative to parametrization of 3-D objects via surfaces. Changes in terms of local translation, bending and widening can be more naturally expressed by medial, rather than surface, representations. Pizer et al. [6, 16] developed an object representation by a mesh of medial atoms. The shape and structure of the skeletal sheet of atoms, as well as a local width function, define the object. Each atom is characterized as a tuple with position, radius, and the normal vectors to the boundary: $m = \{\mathbf{x}, r, \mathbf{n_0}, \mathbf{n_1}\} \in \mathcal{M}$, with $\mathcal{M} = \mathbb{R}^3 \times \mathbb{R}^+ \times \mathcal{S}^2 \times \mathcal{S}^2$. The object surface can be interpolated from endpoints of the sets of medial atoms, but this representation also allows a continuous interpolation of the whole object interior and a rim exterior to the object boundary. Since the parameter vector of medial atoms includes position, length and angle (between normals), mean and variability of a population of object shapes is calculated via the Fréchet mean and PGA framework as discussed before.

Anatomical structures of interest, including left and right hippocampus, amygdala, putamen, caudate, and globus pallidus, have been segmented by well-trained[1] experts using semi-automated procedures. The segmented objects are represented as binary voxel objects (see Fig. 1). We constructed sampled medial models from populations of objects, using the modeling scheme developed by Styner et al. [17] to determine the minimum sampling of each medial mesh model. The m-rep models are deformed to optimally fit the original segmentations of each anatomical object [16]. This process is applied individually to each of the 10 anatomical objects in each of the 20 image datasets. Correspondence across datasets is established by deformable m-rep modeling.

Normalization and Statistics of Pose: The normalization of the m-rep pose is based on a procedure similar to Procrustes analysis. In the m-rep pose normalization,

---

[1] See http://www.psychiatry.unc.edu/autismresearch/mri/roiprotocols.htm .

the sum-of-squared geodesic distances, instead of Euclidean distances, between corresponding medial atoms is minimized, as described in [14]. In this paper we discuss two types of pose normalization in the context of multiple object sets. The first one applies the above procedure to all objects jointly. We call this the global pose normalization of the objects. After this global pose normalization, the individual objects will likely have residual pose variation relative to the global pose (see Fig. 2a,b). Thus, we also perform an object-specific normalization called local pose normalization.

The resulting global-to-local pose change parameters are the features we use to create new statistics of the object pose with a novel PGA extension. As with atoms, the PGA is approximated by mapping the pose parameters to the tangent space and running linear PCA. The mapping is accomplished by taking the log of rotations and scales. The log map of a unit-length quaternion $q = (w, \underline{\mathbf{v}})$ is defined as: $\frac{\theta}{sin(\theta/2)} \cdot \underline{\mathbf{v}}$, where $\theta = 2arccos(w)$. The extension of PGA beyond describing statistics of medial atoms into pose changes allows us to analyze pose and shape simultaneously. This is done by concatenating the pose parameters to the feature vectors of the atoms. Due to their greater magnitude, the pose parameters will tend to dominate the PGA. Therefore, a prewhitening is done to make each feature have standard deviation 1.0 across all samples in the tangent space.

3   Results

Motivation and Clinical Data: The driving clinical problem of this research is the need for a joint analysis of the set of subcortical brain structures, over and above that of individual structures. The image data used in this paper is taken from an ongoing clinical longitudinal pediatric autism study. This study includes autistic subjects (AUT) and typically developing, healthy controls (TYP) with baseline at age 2 and follow-up at age 4. Through this longitudinal design, we can study growth (see Fig. 1), cross-sectional differences between groups and even group growth patterns. For the preliminary analysis shown here, we have selected 5 subjects each from the TYP and AUT groups. For eight of these subjects, we had longitudinal data with successful scans at ages 2 and 4. Our main goal is to systematically study the effect of different pose normalization settings on the analysis of longitudinal and cross-sectional shape changes.

Principle Geodesic Analysis: PGA performs a compression of the multi-object shape variability to a small set of major eigenmodes of deformation. We assume that the first few modes describe most of the shape variability, the rest representing mostly noise. The quality of this compression can be evaluated with the criteria compactness, sensitivity and specificity as discussed in [17]. As a preliminary test, we followed the standard procedure of projecting the multi-object sets into the shape space of the eigenmodes λ_i. This leads to a set of weights in the shape space that describe the deviation of individual shapes from the mean. In our case, each weight vector represents a multi-object shape set. We applied PGA to the whole set of objects for the four subject groups, which ensures projection into the same geometric domain for all subjects.

Shape: Differences in Normalization: First, the sets of objects are aligned purely by a global process, including global translation, rotation and scaling (see Fig. 2b). The resulting PGA captures variability in shape and in residual local pose.
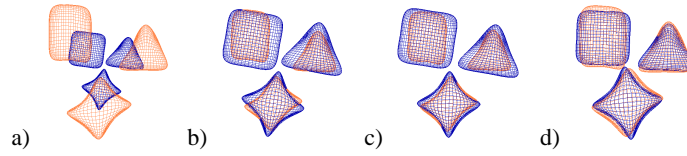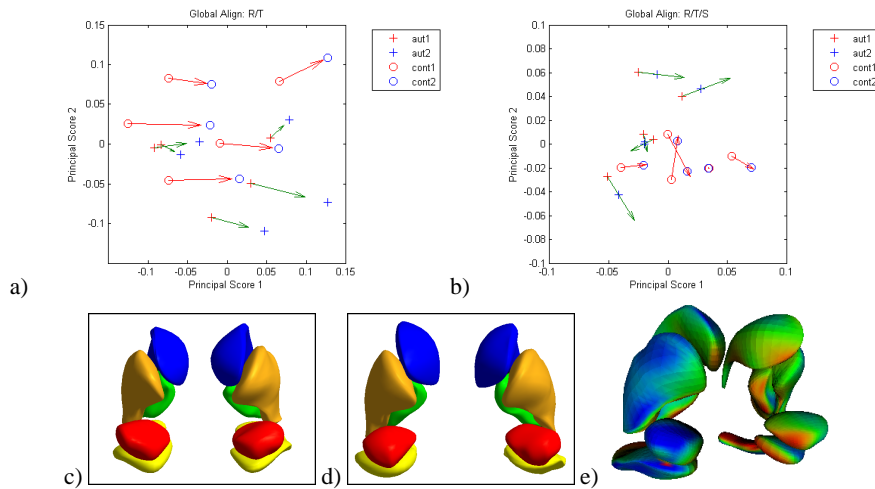
Fig. 2. Normalization examples: a) Global rotation and translation (R/T) normalization; b) global R/T/S; c) global R/T/S plus local R/T normalization; d) global R/T plus local R/T/S

In a second step, we varied the global normalization procedure by disabling scaling normalization and again applied PGA. All objects were thus left in their original size and only rotation and translation were normalized (see Fig. 2a). In order to compare the differences in the two global PGA's (with and without scaling normalization), we plotted the values of the first two major eigenmodes (λ_1, λ_2) of deformation (see Fig. 3a,b). The arrows connect corresponding longitudinal pairs, which allows qualitative evaluation of the correlation between PGA modes and longitudinal changes.

The λ_1 axis in the PGA without global scaling normalization seems to characterize mainly differences between age 2 and 4 (Fig. 3a), as indicated by the parallel alignment of the connecting arrows to the λ_1 axis. After scaling normalization (Fig. 3b) no coherent alignment of the arrows is visible. This suggests that the main effect of longitudinal change is reflected in the scaling normalization and thus the overall size of the object sets. Also, corresponding longitudinal pairs cluster quite well in the plot including scaling normalization, which supports our hypothesis that shape changes due to growth are considerably smaller than shape differences between subjects (see Fig. 1).

As seen in Fig. 2, the individual objects still have relative pose differences after global pose normalization. Therefore, we proceeded by calculating the PGA shape space of objects after a global, then local pose normalization. The top row of Fig. 4 shows the λ_1 vs. λ_2 plot in the global plus local normalization settings. The left plot shows global normalization followed by local normalization including scaling. It is noteworthy that the use of scaling in the global normalization is irrelevant, as the local scaling operation supersedes the global one. This pose normalization setting is similar to the one commonly used in shape analysis of single objects with full Procrustes alignment. The top right plot of Fig. 4 shows the objects with global scaling normalization but no local scaling normalization, which is similar to another common technique of scaling normalization with brain volume.

The high degree of difference between the PGA values in the two different local scaling settings can be due to two different factors: a) inter-subject variability of the residual scaling factors after global normalization or b) instability in the computation of the PGA directions. We are currently working in evaluating the stability of the PGA, but our earlier studies indicate that this cannot be the sole reason for the discrepancy. Also, the arrangement of the groups look quite different in the two plots. In neither plot does there seem to be a clustering according to age, but clustering according to diagnosis is different. Although it is premature to draw any conclusions due to the small sample size

Fig. 3. Top row: λ_1 vs λ_2 plot in pure global normalization settings. a) Global rotation and translation normalization; b) Global R/T/S normalization. The arrows indicate corresponding subject pairs. The range of (a) is larger than (b). Bottom row: c) M-reps deformed to -3 standard deviations of the first eigenmode of PGA with global R/T; d) +3 standard deviations; e) Euclidean distance map between mean of global R/T aligned m-reps at age 2 and age 4. The colormap ranges from red (Age 2 > Age 4) over green (Age 2 = Age 4) to blue (Age 2 < Age 4).

and limited knowledge about the PGA stability, the preliminary results demonstrate the potential of multi-object shape analysis.

Pose and Shape: Local Pose Change Parameters: The purpose of the normalization discussed previously is to decouple pose and shape differences in order to study only shape in the PGA. In a multi-object setting, however, the pose changes may be of interest. Therefore, we ran PGA on the local translation, rotation, and scale pose parameters as seen in Fig. 4, lower left. These represent the local residual pose changes after a global translation and rotation alignment. An advantage of this PGA calculation is that it suffers less from the high dimensionality, low sample size problem of the PGA on the medial atoms. Here, each subject is represented by a feature vector of length 70 (10 objects, each with translation, rotation, scale) as opposed to about 1900 for the medial atoms (210 atoms, each with 9 parameters). By analyzing only pose, we see a somewhat similar separation according to diagnosis as in the upper right of Fig. 4.

Having analyzed pose and shape separately, our final step was to run PGA on both simultaneously. The lower right plot of Fig. 4 shows that the prewhitening discussed earlier has a large effect on the eigenmodes as compared with the upper left. A prewhitened PGA on only the medial atoms is not altered significantly by including the pose parameters, meaning the prewhitening itself accounts for the changes between the two plots and that including the pose does not destabilize the computation.
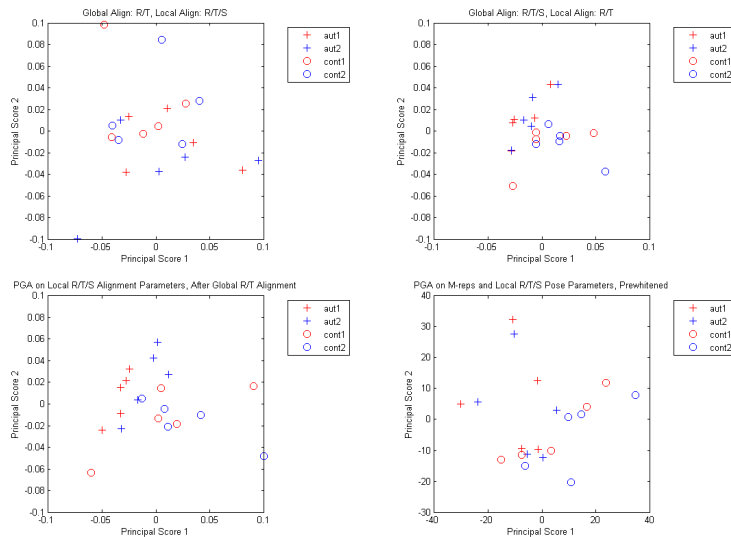
Fig. 4. λ_1 vs λ_2 plots of different PGA's. Top row: The left plot shows global rotation and translation normalization followed by local rotation, translation and scaling. The right plot has global scaling normalization, but no local scaling normalization. Bottom row: The left plot includes only the local pose change parameters. The right uses both medial atoms and pose parameters.

4 Discussion

We have discussed work in progress towards extending statistical analysis of anatomical shape from single structures to multi-object sets. Key issues addressed are the extraction of a small set of key features representing both the shape and pose of object sets and calculation of mean and variability via Riemannian metrics. The current results suggest that after removing global scale, longitudinal data of the same subject cluster closely in the PGA space and thus that longitudinal shape change is smaller than shape variability across subjects, a driving hypothesis shown in Fig. 1. Also, projections of subject groups into PGA components vary significantly depending on choice of normalization.

Several open issues remain and need to be addressed by our group and the international research community. In regard to the m-rep object parametrization used here, we still need to demonstrate the quality and stability of correspondence, as well as the robustness, sensitivity and specificity of PGA-based compression of features. Although our results indicate a possible separation between groups, it is important to note that PGA, similar to PCA, selects a subspace based on maximum common variability, not maximum separation. An extension of independent component analysis (ICA) to curved space or supervised training of a subspace of maximum separation will be explored in our future research. Applications in neuroimaging further require hypothesis testing schemes that combine shape and pose features with clinical variables, and that have to properly address the problems of nonlinear modeling and multiple comparison testing. Encouraging progress is shown by recent work of Terriberry et al. [18].

References

1. Dryden, I., Mardia, K.: Multivariate shape analysis. Sankhya **55** (1993) 460–480
2. Small, C.G.: The statistical theory of shape. Springer (1996)
3. Csernansky, J., Joshi, S., Wang, L., Haller, J., Gado, M., Miller, J., Grenander, U., Miller, M.: Hippocampal morphometry in schizophrenia via high dimensional brain mapping. Proc. Natl. Acad. Sci. USA **95** (1998) 11406–11411
4. Tsai, A., Yezzi, A., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, E., Willsky, A.: Shape-based approach to curve evolution for segmentation of medical imagery. IEEE Transactions on Medical Imaging **22**(2) (2003) 137–154
5. Yang, J., Staib, L.H., Duncan, J.S.: Neighbor-constrained segmentation with level set based 3d deformable models. IEEE Transactions on Medical Imaging **23**(8) (2004) 940–948
6. Pizer, S., Fritsch, D., Yushkevich, P., Johnson, V., Chaney, E.: Segmentation, registration, and measurement of shape variation via image object shape. In: IEEE Trans. Med. Imaging. Volume 18. (1999) 851–865
7. Bookstein, F.: Shape and the information in medical images: A decade of the morphometric synthesis. In: MMBIA. (1996)
8. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models - their training and application. In: Computer Vision and Image Understanding. (1995) 38–59
9. Kelemen, A., Székely, G., Gerig, G.: Elastic model-based segmentation of 3d neuroradiological data sets. In: IEEE Trans. Med. Imaging. Volume 18. (1999) 828–839
10. Davatzikos, C., Vaillant, M., Resnick, S., Prince, J., Letovsky, S., Bryan, R.: A computerized method for morphological analysis of the corpus callosum. J. of Comp. Assisted Tomography **20** (1996) 88–97
11. Thompson, P., Mega, M., Toga, A.: Disease-Specific Brain Atlases. In: Brain Mapping: The Disorders. Academic Press (2000)
12. Thompson, P., Giedd, J., Woods, R., MacDonald, D., Evans, A., Toga, A.: Growth patterns in the developing brain detected by using continuum mechanical tensor maps. Nature **404** (2000) 190–193
13. Davis, B., Lorenzen, P., Joshi, S.: Large deformation minimum mean squared error template estimation for computational anatomy. In: ISBI. (2004) 173–176
14. Fletcher, P., Lu, C., Pizer, S., Joshi, S.: Principal geodesic analysis for the study of nonlinear statistics of shape. In: IEEE Transactions on Medical Imaging. Volume 23. (2004) 995–1005
15. Fletcher, P.: Statistical Variability in Nonlinear Spaces: Application to Shape Analysis and DT-MRI. PhD thesis, The University of North Carolina at Chapel Hill (2004)
16. Pizer, S., Fletcher, T., Fridman, Y., Fritsch, D., Gash, A., Glotzer, J., Joshi, S., Thall, A., Tracton, G., Yushkevich, P., Chaney, E.: Deformable m-reps for 3d medical image segmentation. International Journal of Computer Vision IJCV **55**(2) (2003) 85–106
17. Styner, M., Gerig, G., Lieberman, J., Jones, D., Weinberger, D.: Statistical shape analysis of neuroanatomical structures based on medial models. Medical Image Analysis (MEDIA) **7**(3) (2003) 207–220
18. Terriberry, T., Joshi, S., Gerig, G.: Hypothesis testing with nonlinear shape models. In Christensen, G.E., Sonka, M., eds.: Information Processing in Medical Imaging (IPMI). Number 3565 in Lecture Notes in Computer Science LNCS, Springer Verlag (2005)