Localized Components Analysis of Shape Variation

Dan Alcantara[1a], Owen Carmichael[1ab], Will Harcourt-Smith[2,3], Kirstin Sterner[3,4], Stephen Frost[3,5], Rebecca Dutton[6], Paul Thompson[6], Nina Amenta[1a]

1: Computer Science(a) and Neurology(b) Departments, University of California, Davis

2: American Museum of Natural History

3: NYCEP Morphometrics Group

4: Anthropology Department, New York University

5: Anthropology Department, University of Oregon

6: Neurology Department and Laboratory of Neuro Imaging, University of California, Los Angeles

Abstract

Localized Components Analysis (LoCA) is a new method for describing surface shape variation in an ensemble of objects using a linear subspace of *spatially-localized* shape components. In contrast to earlier methods, LoCA optimizes explicitly for localized components and allows a flexible trade-off between localized and concise representations, and the formulation of locality is flexible enough to incorporate properties such as symmetry. We demonstrate the application of LoCA to the detection of shape differences associated with sex, disease state, and species.

I. Introduction

The parameterization of an ensemble of shapes is a key step in a broad array of applications that require quantification or manipulation of the shape properties of objects. In this paper, shape parameterization refers to the problem of converting a representation of the delineating boundary of an object in 2D or 3D into a concise vector of numbers that captures its salient shape characteristics. Converting the shape of an object such as an organ or bone into a small set of *shape parameters* facilitates a variety of statistical analyses, including the characterization of shape variability across an ensemble; comparison of object shape between groups; and the tracking of shape change over time.

Often, converting shapes into concise parameter vectors is intended to make the compression, transmission, classification, and modification of shape data more computationally efficient. In contrast, we focus on finding an intuitive shape parameterization. In biology and medicine, presenting the results of shape analyses in an intuitive way can encourage the connection of shape properties to domain-specific physical or biological processes; for instance, the interpretability of brain shape parameterization could be enhanced if each parameter represents an easily-grasped aspect of the brain, such as the size of an identifiable, anatomically-localized brain region. The use of an intuitive parameterization could in turn promote interpretations of brain shape differences between healthy and sick individuals in terms of disease causes or effects. In graphics and vision, using each parameter to represent a single, intuitive aspect of shape could enhance the usability of user interfaces for shape model manipulation and could simplify automated methods for object recognition based on the shape of object parts. However, we note that for some applications, simultaneously representing distal, correlated shape characteristics in the same shape parameter may be desirable for shape compression or analysis.

Here, we define an intuitive paramaterization to be one that is both *concise* (capturing salient shape characteristics in a small number of parameters) and *spatially localized* (accounting for the shape of a spatially restricted sub-region in each parameter). Our hypothesis is that spatially-localized shape parameterizations are more intuitive, since they allow users to conceptualize object shape in terms a small number of object parts, which might be affected differentially by physical phenomena, or may simply be of interest. In the above brain shape example, for instance, shape change due to disease processes is known to occur in spatially-localized brain sub-regions in a variety of disorders (*e.g.*, [5]). Concise parameterizations are attractive both for comprehensibility and because the statistical power of tests on those parameters is reduced as little as possible by corrections for multiple comparisons [8]. While characterizing the deformation of localized regions is not the goai of every analysis of shape variation – sometimes the focus is on the correlation of deformations of distal characteristics, for instance – it is often useful.

We follow the *linear subspace* paradigm of expressing each shape as a linear combination of prototypical, or *basis* shapes. That is, if each shape is represented as a vector $\mathbf{v}_j$ of the $2m$ or $3m$ coordinates of m points sampled from its boundary (*i.e.*, $\mathbf{v}_j = [\mathbf{v}_{j,1}, \mathbf{v}_{j,2}, \cdots \mathbf{v}_{j,m}], \mathbf{v}_{j,k} = [x_k, y_k]$ for 2D shapes), $\mathbf{v}_j$ is approximated as a linear combination of k basis vectors $\{\mathbf{e}_1, \mathbf{e}_2, \cdots \mathbf{e}_k\}$:

$$\mathbf{v}_j^k = \sum_{i=1}^{k} \alpha_{j,i} * \mathbf{e}_i$$

The shape parameters are the coefficients $\alpha_{j,i}$. Linear subspace methods are attractive because their linearity in $\mathbf{e}_i$ allows them to be manipulated using standard tools from linear algebra.

Figure 1 (left) depicts a typical $\mathbf{e}_i$ generated by the classical linear subspace method, Principal Components Analysis (PCA), applied to tracings of human *corpora callosa* (CC). The basis shape summarizes a complex pattern of shape characteristics across the entirety of the CC. Therefore, if the corresponding α_i differs between groups, the explanation of the group difference in physical terms is complex. Localized Components Analysis (LoCA), in contrast, optimizes the $\mathbf{e}_i$ for spatial locality and conciseness simultaneously. It improves on previous linear subspace methods by explicitly optimizing for localized shape parameters and by allowing the user to modulate the trade-off between locality and conciseness with greater flexibility than previous methods. Figure 1 (right) shows a vector from a completely localized LoCA basis; differences in the corresponding α_i between groups gives rise to a simple physical explanation in terms of

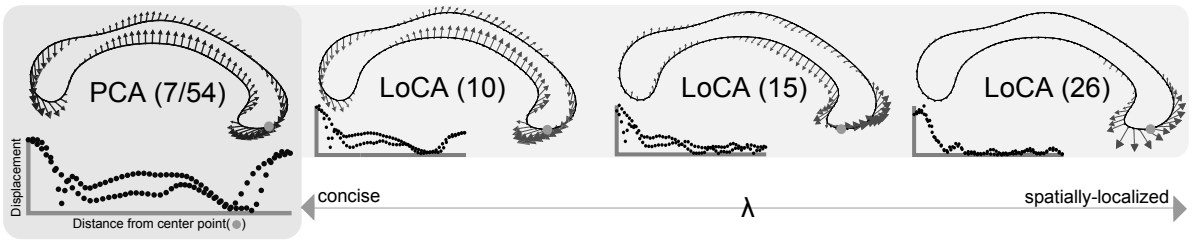

Fig. 1. The vector capturing the most shape variation in the corpora callosa data set is shown for PCA and LoCA bases with increasing spatial locality. Arrows show how points are affected as the corresponding shape parameter is varied. The PCA basis and most concise LoCA basis are equivalent, capturing 90% of the shape variation with fewer vectors (in parentheses), but consisting of vectors representing complex, global patterns of shape characteristics. The accompanying graph shows that points distant from the *center point* (see Section III) are heavily displaced. As conciseness is traded in favor of locality, the LoCA vectors focus on one particular characteristic (the expansion of the genu).

the *genu*, the CC subregion whose shape is captured by the $\mathbf{e}_i$.

LoCA was introduced in a related conference paper [1]. That paper gave in-depth experimental examples of LoCA on three data sets, comparing the locality and conciseness of LoCA to PCA and related techniques. In this paper, we begin by summarizing related techniques in Section II, and then we present LoCA in Section III (for more details and examples see [1]). Then in Section IV, we use LoCA to illustrate the shape differences between groups in three datasets. We examine how well-established sex differences in two-dimensional sections of human CCs are expressed by LoCA, we compare LoCA with a complementary technique for the visualization of shape changes in the lateral ventricles of human brains associated with HIV/AIDS, and we use LoCA to isolate specific shape differences in the crania of closely-related colobine monkeys.

II. Related Work

PCA has been used to find concise bases for shape spaces in medical image analysis [7], morphometrics [4], computer graphics [2], and many other contexts. In PCA, $\mathbf{e}_i$ is the ith eigenvector of the covariance matrix of the example $\mathbf{v}_j$ vectors; therefore, the $\mathbf{e}_i$ are orthogonal and $\mathbf{v}_j^k$ is the best k-th order approximation of $\mathbf{v}_j$ under the L_2 norm. Related algorithms such as Sparse PCA (S-PCA) [6] [23] [15], independent components analysis (ICA) or principal factor analysis (PFA) do not directly optimize a locality-related objective function, but they appear to generate spatially-localized components anyway [15] [16] [20] [3]. Alternatively, pre-defined spatially located regions of interest can be integrated into PCA [21]. Our approach is inspired

by S-PCA and follows a similar strategy of adjusting the $\mathbf{e}_i$ provided by PCA; but we explicitly optimize for locality and allow the user to explicitly modulate the trade-off between conciseness and locality.

Other recent techniques have been proposed to capture localized shape differences. Networks of localized medial primitives [13] seem computationally more difficult than LoCA, while direct comparison of corresponding $\mathbf{v}_{j,k}$ [5] suffers from a reduction in statistical sensitivity as compared to methods employing dimension reduction. We compare LoCA directly with the latter technique in Section IV.

III. METHODS

PCA produces the most concise basis possible under the L_2 norm; that is, for each k, $\sum_{j=1}^{n} ||\mathbf{v}_j - \mathbf{v}_j^k||_{L_2}$ is minimized when $\mathbf{e}_1 \cdots \mathbf{e}_k$ are the first k eigenvectors of the covariance matrix of the $\mathbf{v}_j$. We use a formulation of PCA as the minimization of an energy function E_{var} as in [6], and modify it by minimizing $E_{var} + \lambda E_{loc}$, where E_{loc} is a new energy term that summarizes the spatial locality of the $\mathbf{e}_i$. The λ balances the trade-off between the competing interests of conciseness and locality (Figure 1).

A. Energy Function

Each successive PCA component accounts for as much shape variation as possible; that is, the distribution of shape variation over the PCA basis vectors is as concentrated as possible on the leading $\mathbf{e}_i$. More formally, one can define the relative variance β_i of each basis vector $\mathbf{e}_i$ as

$$\beta_i = \frac{\sum_{j=1}^{n} \langle (\mathbf{v}_j - \mu), e_i \rangle^2}{\sum_{j=1}^{n} ||\mathbf{v}_j - \mu||^2}$$

where μ represents the mean of the data vectors $\mathbf{v}_j$. The entropy of the distribution of relative variances, $-\sum_{i=1}^{k} \beta_i \log \beta_i$, is minimized, over all orthogonal bases, by the PCA basis, so we define this to be E_{var}, as in [6]. Each $\mathbf{e}_i$ consists of a concatenation of sub-vectors $\mathbf{e}_{i,j}$, each of which correspond to the spatial coordinates of one of the surface points (*i.e.*, $\mathbf{e}_i = [\mathbf{e}_{i,1}, \mathbf{e}_{i,2}, \cdots \mathbf{e}_{i,m}], \mathbf{e}_{i,j} = [\mathbf{e}_{i,j,x}, \mathbf{e}_{i,j,y}]$ for 2-D shapes). We encourage each $\mathbf{e}_i$ to have simultaneous nonzero entries corresponding to points p_1 and p_2 (*i.e.*, simultaneous nonzero entries in $\mathbf{e}_{i,1}$ and $\mathbf{e}_{i,2}$) if and only if p_1 and p_2 are close to each other. To do so, we introduce a *pairwise compatibility matrix* $\mathbf{B}$ whose entries $\mathbf{B}[i,j]$ tend toward 1 when p_i and p_j are near each other,

and tend towards 0 when they are distant. In a variant of our method, called Symmetric LoCA, $\mathbf{B}[i,j]$ tends toward 1 when p_i and p_j are either close or symmetric: near each other after one of the points is reflected across a symmetry plane defined on the object. The **B** matricies we use for our experiments vary as appropriate for the dataset. We use the **B** matrix to define a cost function $C(\mathbf{e}_i, p_c)$, which evaluates each surface point p_c as a potential *center point* for $\mathbf{e}_i$. The surface point p_c is a good center point for $\mathbf{e}_i$ if the deformation described by e_i is localized around p_c. Specifically:

$$C(\mathbf{e}_i, p_c) = \sum_{j=1}^{m} |\mathbf{B}[c,j] - ||\mathbf{e}_{i,j}||_{L_2}|^{\kappa}$$

The $\mathbf{e}_i$ have unit length, so both $\mathbf{B}[c,j]$ and $||\mathbf{e}_{i,j}||$ vary between 0 and 1. Center points p_c that $\mathbf{e}_i$ move significantly, and for whom the entries in $\mathbf{e}_i$ respect the compatibilities between p_c and other surface points, incur a low cost. We found experimentally that setting κ to a value less than 2 (1.5 in all experiments below) helped us ameliorate outlier effects in our experiments.

A vector $\mathbf{e}_i$ is well-localized if we can find one (or more) good center points p_c, but we need to consider one additional complication: some p_c may be favored as center points by the object geometry, as expressed by the compatibilities in matrix **B**, irrespective of the deformations in the input data $\mathbf{v}_j$ (specifically, those near the center of the object). Therefore, for a given p_c, we normalize $C(\mathbf{e}_i, p_c)$ by the maximum value that $C(\mathbf{e}, p_c)$ could take over all possible **e**.

$$E_{loc} = \sum_i \min_{p_c} \frac{C(\mathbf{e}_i, p_c)}{\max_{\mathbf{e}} C(\mathbf{e}, p_c)}$$

The denominator for a given p_c is simply $\sum_j max(|B[c,j]-1|, |B[c,j]-0|)^{\kappa}$. It needs to be computed only once.

B. Optimization

We assume that we are given an ensemble of n objects, each represented by m points on its boundary, and the compatibility matrix **B**. Overall differences in object scale, rotation and translation over the ensemble are removed through generalized Procrustes alignment [4]. The resulting scaled and aligned data sets are used as input to the numerical optimization.

Our optimization procedure is similar to that used in [6]. PCA provides an initial orthonormal basis **e**, and every possible pair $\mathbf{e}_i, \mathbf{e}_j$ are rotated together in the two-dimensional plane they span, so that the basis remains orthonormal. Each pair is rotated by the angle θ that minimizes

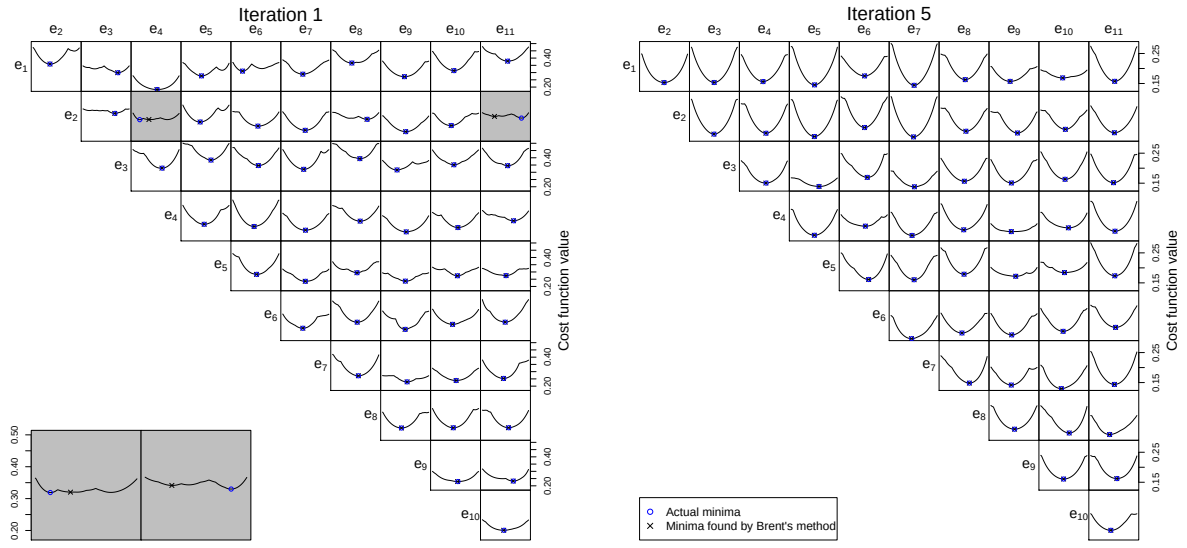

Fig. 2. Behavior of the cost function when rotating pairs of the 10 basis vectors accounting for the most variation. These graphs, generated from the ventricle data set, are representative of the graphs for the other data sets. In almost all cases, the global minimum is found by Brent's method. If local minima are found, it is typically due to a flat cost function where many rotation angles provide nearly the same reduction in cost.

$E_{var} + \lambda E_{loc}$. The optimal θ is found numerically using Brent's method [14]. As illustrated in Figure 2, the global minimum is almost always found at each iteration, even though changes of center point produce discontinuities.

The pairs are rotated in decreasing order of shape variation accounted for. The set of all $\mathbf{e}_i, \mathbf{e}_j$ pairs are adjusted repeatedly, and optimization ceases when adjusting them changes the objective function less than a fixed threshold. Between 50 and 150 iterations were required for each experiment below.

IV. Results

Below, we compare LoCA to PCA on three data sets, two from brain imaging and one from primate morphology: CCs, lateral ventricles, and colobine monkey skulls[1].

LoCA behavior depends strongly on λ, the parameter that modulates the trade-off between conciseness and locality (Figure 1). For $\lambda = 0$, LoCA reduces to PCA. For small λ, LoCA basis

[1]Movies and larger images are at: http://idav.ucdavis.edu/∼dfalcant/loca.html

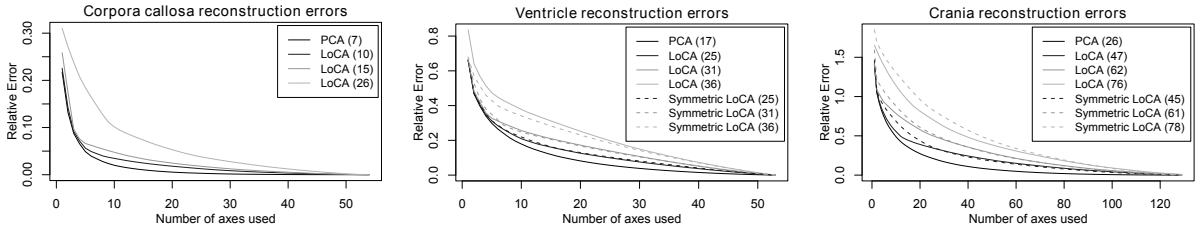

Fig. 3. Reconstruction error when using only the first k vectors of the basis. The numbers in parentheses denote the number of vectors required to capture 90% of the shape variation. For the corpora callosa, these numbers correspond to the different λ settings used in Figure 1. When available, LoCA bases are compared with Symmetric LoCA bases requiring a similar number of vectors. Symmetric LoCA provides a small tradeoff between reconstruction error and symmetry across the midsagittal plane.

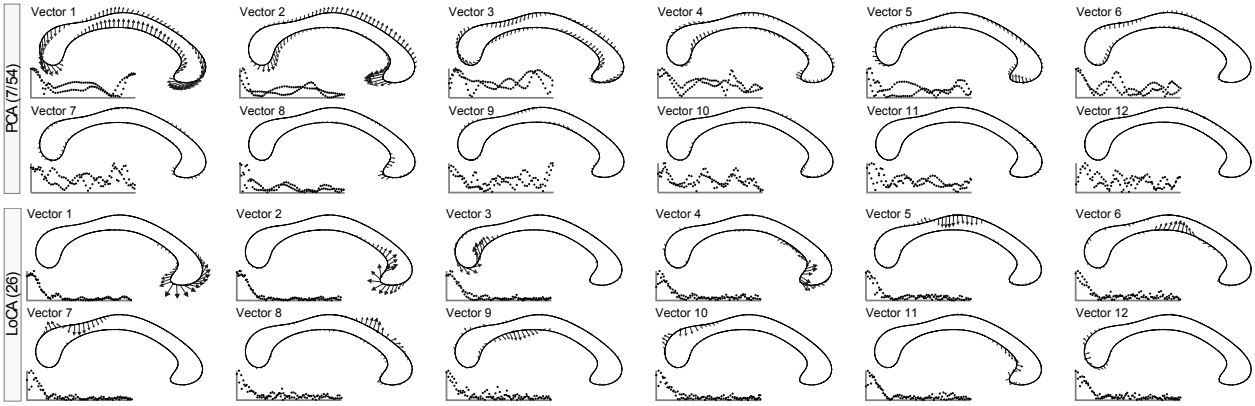

Fig. 4. Corpora callosa basis comparison. The first 12 out of 54 basis vectors are shown. LoCA successfully captures the major shape deformations of the genu and splenium in the first four vectors.

vectors accounting for the highest amounts of shape variation resemble PCA basis vectors, while the rest of the basis is clearly localized. For larger λ, all LoCA basis vectors are local, and the bases require more basis vectors to account for shape variation in the data. In Figures 4, 5, and 8, LoCA basis vectors are depicted for the smallest value of λ for which the bases lacked global basis vectors. Reconstruction error for all three datasets is graphed in Figure 3, in which we can see that strongly local bases (*i.e.*, higher λ) require more vectors for accurate reconstruction.

A. Corpora callosa

The CC data set represents 31 AIDS patients and 19 HIV-seronegative controls (mean age 42.64 +/- 11.28 S.D.), who underwent high-resolution magnetic resonance brain scans as part

of a previously-described study [18]. As part of that study, each scan was rigidly aligned to the ICBM-53 template based on intensity cues, and all scans were then further aligned to each other in a groupwise fashion, based on manual sulcal landmarks. The CC was manually traced on the midsagittal slice of each scan using a reliable and repeatable protocol, partitioned using the Witelson criteria [22], and automatically marked with 103 surface correspondences [10].

We compared the ability of LoCA and PCA to detect what is likely the most firmly-established difference in CC morphology between men and women, based on post-mortem [22] and imaging-based [11] [17] findings: that women have larger isthmuses (*i.e.* posterior midbodies). We selected the LoCA and PCA coefficients for the shape components that captured 90% of the CC shape variation: the top 7 for PCA, and the top 26 for LoCA (see Figure 4). For each shape component, gender differences in its coefficients were tested in linear statistical models that controlled for other factors known to affect CC morphology: presence of AIDS, handedness, and age. The significance of the gender differences were assessed by F tests in an ANOVA design.

Among the top 26 LoCA shape components, only components 5 and 9 had coefficients that differed significantly by gender at the $p < .05$ level, indicating thicker superior and inferior aspects of the isthmus in women (component 5: $F = 5.44$, $p = 0.024$; component 9: $F = 4.24$, $p = 0.045$, p values uncorrected). This gender difference has been reported extensively (***e.g.***, [22] [11] [17]). Only the first PCA component differed significantly between men and women ($F = 5.41$, $p = 0.024$); it represents a complex pattern of elongation, rotation, bending and twisting of the entire structure. The complexity makes the shape component relatively difficult to interpret in simple anatomical terms. Also, some of the represented modes of CC deformation have not been reported to differ between genders in the literature, so their validity is uncertain.

B. Lateral Ventricles

Lateral ventricles were also traced manually in 3D on 54 of the 55 group-aligned scans of healthy and HIV/AIDS patients from a prior study [18]. The boundaries of the frontal, temporal and occipital horns were traced on each slice based on a reliable, repeatable protocol described previously [12]. Dense one-to-one correspondences between subjects at homologous surface points were established by threading medial curves down the center of each of the frontal, temporal, and occipital horns, and resampling each horn to contain a fixed number of axis-aligned parallel traces [18]. Within each trace, rays were cast outward from the medial curve point toward

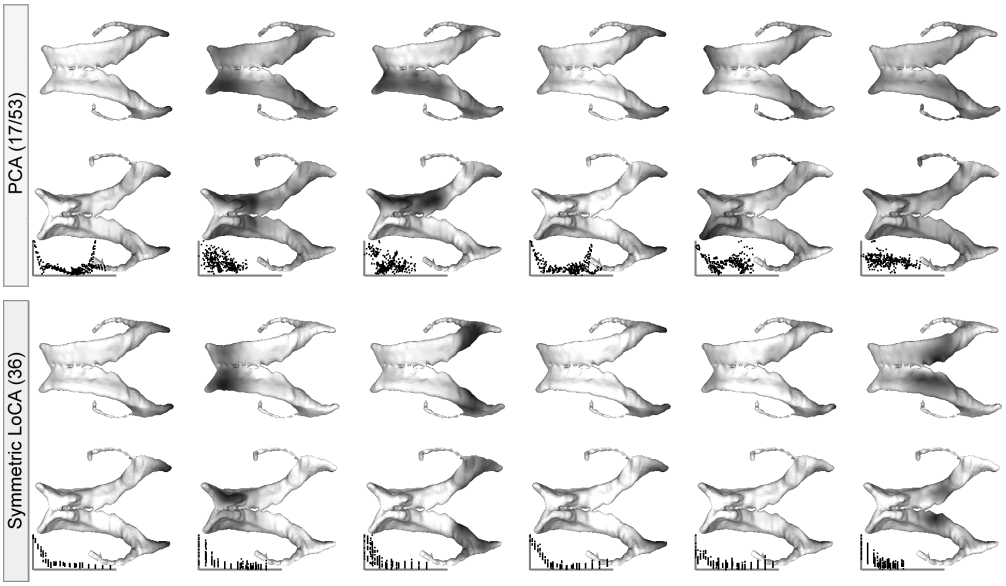

Fig. 5. Top and bottom views of the first 6 basis vectors from the ventricle data set, ordered by reconstruction error and colored by displacement magnitude. Symmetric LoCA affects corresponding parts of the ventricles at the same time: its first vector lengthens both occipital horns, while its second expands both frontal horns.

the ventricular surface every θ radians. Across subjects, 330 correspondences were established between points in analogous traces and analogous θ. Geodesic distances were computed between pairs of resampled surface points, and the distances were used to compute compatibilities. The resampled surface points were the input to the shape parameterization methods.

We used PCA and Symmetric LoCA to assess differences in shape component coefficients between AIDS and control groups, and compared the results to differences between these groups that were previously presented using an established radial ventricular mapping approach [18]. The radial mapping approach has been strongly validated and used to characterize ventricular shape variation across a large number of medical covariates [12][5][19]. At each surface point, local thickness values were computed by calculating the distance from the surface point to the medial curve; mean thickness values were computed for AIDS and control groups. The ratios of mean local thicknesses between AIDS and control groups were computed and mapped to color values in a 3D rendering of the ventricular surface (Figure 6). Red regions represent portions of the ventricular surface that were extremely dilated in AIDS patients compared to controls. Coefficients for the top PCA and LoCA shape components accounting for 90% of ventricular shape variation were analyzed in statistical models (17 PCA components and 36 LoCA

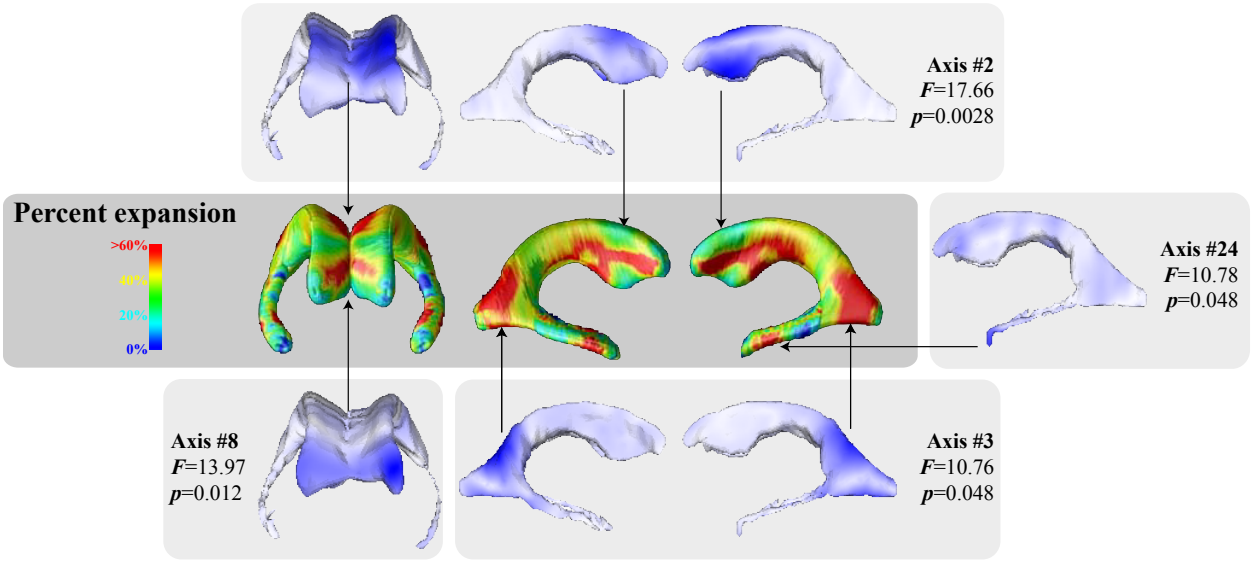

Fig. 6. Differences between AIDS and control groups in ventricular shape as characterized by LoCA and a complementary radial thickness mapping approach. Regions that are extremely dilated in the AIDS group are shown in red in the spatial map; LoCA shape components whose coefficients differed significantly between AIDS and control groups are also shown to illustrate the agreement between the methods. Shown p values have been Bonferroni-corrected to correct for multiple comparisons. The spatial map was adapted from a similar figure in [18].

components; the initial vectors in each basis appear in Figure 5). The LoCA shape components that differed significantly at the $p < .05$ level between groups, in two-tailed t tests with Bonferroni correction for multiple comparisons, are shown in Figure 6. The ventricular regions covered by these shape components roughly correspond to the (red) regions that are grossly dilated in AIDS patients on the spatial maps, including the lateral walls, superior ridge and anterior tip of the frontal horns; the superior and lateral aspects of the occipital horns; and the tips of the temporal horns. The only statistically significant PCA component was the second (Figure 5, $p = 8.59e{-}05$, $t = 26.25$), which represents nearly uniform inflation of the entire ventricular surface with no indication that some portions of the surface are more or less dilated in the AIDS group.

C. Colobine monkey crania

A shape space was built from a set of 238 colobine monkey crania (Subfamily Colobinae, Family Cercopithecidae) representing seventeen species (two of them outgroups, not pictured in Figure 9). Various primate morphologists collected 45 landmark points on each cranium using a

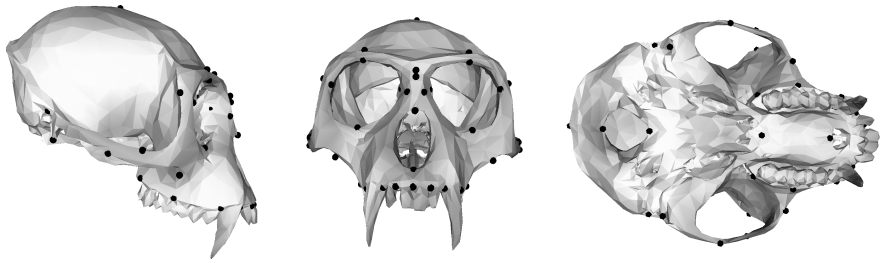

Fig. 7. Locations of the 45 landmark points, shown on the *Colobus guereza* cranium, warped to match the average landmark locations over the whole dataset.

Microscribe 3D digitizer as part of data collection for a long-term project on Old World monkey cranial evolution. Twenty-nine of the landmarks were collected from the dorsal aspect of the cranium, and sixteen were collected from the ventral aspect (Figure 7, and see [9] for details). The landmark sets were rescaled to remove the influence of size differences, and aligned using GPA. Since we did not have surface models for each of the specimens, but only the landmarks, we estimated geodesic distance between landmarks by connecting them with a graph, weighted by distance, representing adjacency on the skull surface, and using distance in the weighted graph. Results of the LoCA and Symmetric LoCA computations are visualized in figure 8.

We used the smallest value of λ which produced a Symmetric LoCA basis which was mostly localized; it required 45 components to account for 90% of the shape variance. We then compared pairs of sister taxa in the tree, testing the statistical significance of the PCA and LoCA vectors for discriminating between the two groups. We employed Welch's two sample t-tests, using Bonferroni correction for multiple comparisons. In each case, we found several statistically significant LoCA components with corrected p-values at or below 0.05. Three examples are shown in Figure 9; we chose these three because there were roughly equal numbers of males and females in the populations sampled, so that the examples can reasonably be interpreted as species and not sex differences, (which are in any case not well-captured by this landmark set).

For LoCA, vector 8 captures one difference between *Presbytis melalophos* and *Presbytis potenziani*. It represents a deformation of the sphenoid bone, internal to the zygomatic arch (cheekbone) and just behind the tooth row. The dramatic difference at the top of the head is also represented in the Symmetric LoCA basis by Vector 19, with a corrected p-value of $5.32e-7$, but we show Vector 8 instead since the brain case is not sampled well by this landmark set.

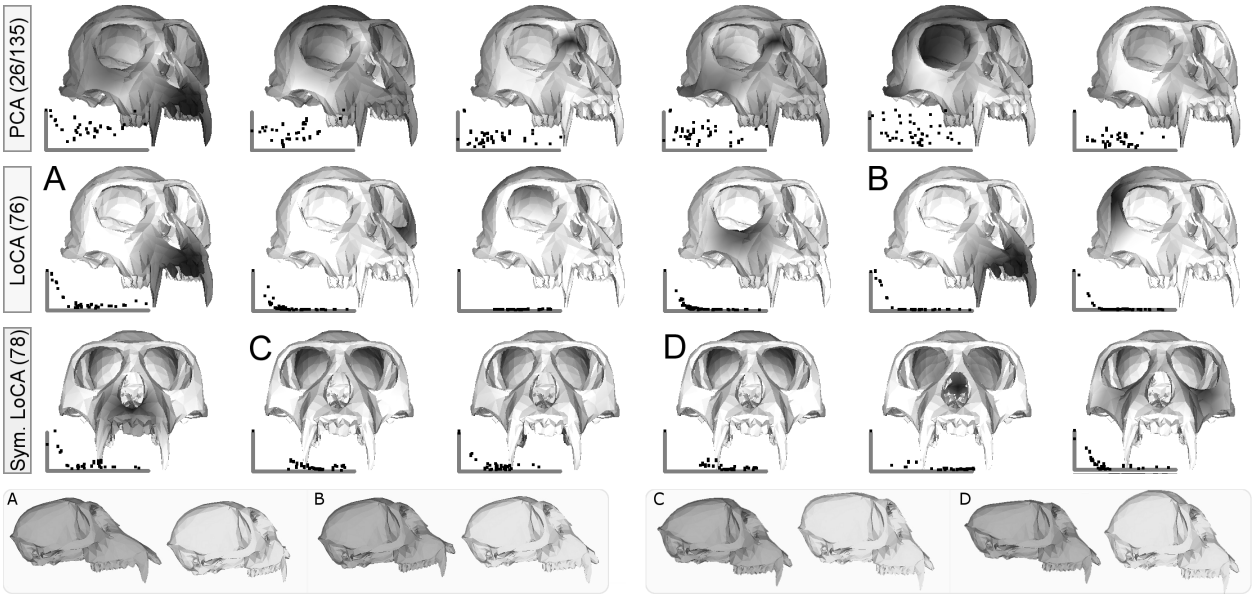

Fig. 8. Cranial basis comparison. For visualization, displacements of the landmark points by each vector are computed and then interpolated onto a *Colobus guereza* cranium using a thin-plate-spline. Darker locations indicate greater displacement magnitudes; because the crania are shown from a single angle, vectors representing motion on the bottom of the cranium appear completely white. *Bottom*: alternative views of similar looking cranial basis vectors. A shows prognathism (snout elongation), while B shows facial kyphosis (teeth straightening). C and D both affect a single landmark point at the top of the skull, with one moving it vertically and the other horizontally.

One of the vectors differentiating between *Nasalis larvatus* and *Simias concolor* is vector 20, which affects the vertical angle of the zygomatic arches. Vector 5 represents a deformation of the nasal bone; it is extended above the snout in *Procolobus verus*, but lays behind it in *Piliocolobus badius*. There are also PCA vectors which discriminate between these three sibling pairs, with low p-values, but they represent complicated deformations which are not easily described, as shown in Figure 9.

V. Discussion

The LoCA components spread the shape variation in the ensemble across more significant components than PCA does, with each component optimized for spatial locality. These components isolate the shape variation of specific parts of the object. Because we often understand shape variation in terms of spatially-localized object parts, this decomposition is useful in the analysis and interpretation of shape differences between groups. Also, we found that representing

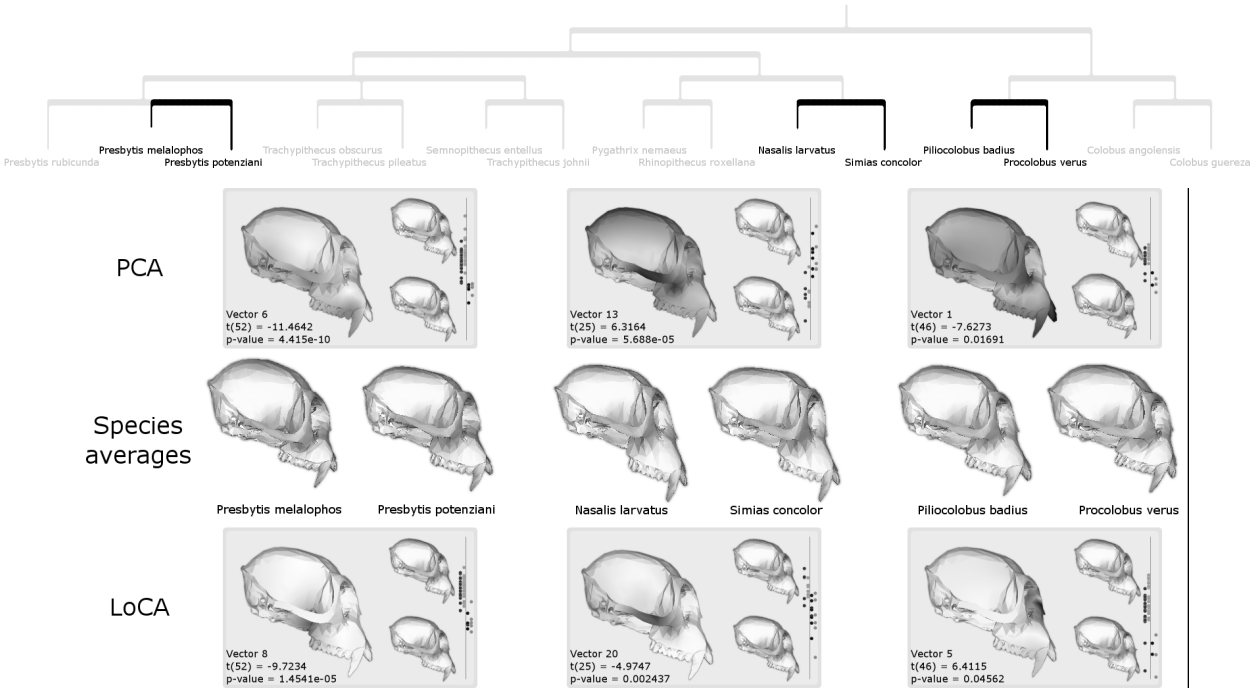

Fig. 9. T-tests comparing the coefficients of three sister taxa pairs in the crania data set were run on PCA and Symmetric LoCA bases. Vectors with small Bonferroni corrected p-values are shown next to crania exbihiting their effects on the average model. To the right of each vector, the distribution of coefficients is shown for the two species (left and right) and genders (light and dark). Each PCA vector shows several shape changes, while the LoCA vectors focus on a single feature that can be intuitively understood.

the shape of single, rather than multiple, object regions in a single shape parameter helped to prevent the obfuscation of relationships between individual regions and auxiliary variables.

We found that LoCA is broadly applicable to a variety of data sets despite differences in dimension, shape representation, and the density, accuracy, and source of the point-to-point correspondences. This flexibility suggests that it might be applied even more broadly, particular to articulated forms. The method is also flexible with respect to the definition of compatibility between surface points; the incorporation of symmetry that we used here is just one example. Another appealing possibility would be somehow combining geodesic distance with curvature change, to encourage vectors affecting regions of similar curvature (eg. the ridges of the ventricles).

In our formulation, pairs of points are more compatible if the geodesic distance between them is lower. Hence we implicitly encouraged a fixed neighborhood size for our localized basis

vectors. It may be more natural to encourage shape basis vectors in which all of the points in a particular surface neighborhood deform in concert, regardless of the size of the neighborhood. Future work will address the problem of how to construct shape bases whose basis vectors are spatially-localized in this scale-invariant fashion.

It is unclear whether we would generate shape parameterizations that are more spatially-localized, or more useful for making shape-based inferences, if we dropped the orthogonality constraint from LoCA or added other constraints to the optimization; for example, we could seek shape components that are both spatially localized and statistically independent, as in ICA. Future work will address this variant of our current method.

VI. Acknowledgments

This research was funded by grants NSF IIS-0513660 and NSF IIS-0513894. We thank Professor Howard Aizenstein, Professor Oscar Lopez, and Professor James Becker for their roles in collecting the CC and ventricle data. We thank Professor Eric Delson for his role in collecting the colobine monkey crania. This is NYCEP Morphometrics contribution number 32.

References

[1] D. Alcantara, O. Carmichael, E. Delson, W. Harcourt-Smith, K. Sterner, S. Frost, R. Dutton, P. Thompson, H. Aizenstein, O. Lopez, J. Becker, and N. Amenta, "Localized components analysis," in *Proc. IPMI*, 2007, pp. 519 – 531.

[2] B. Allen, B. Curless, and Z. Popovic, "The space of human body shapes: reconstruction and parameterization from range scans," in *Proceedings of ACM SIGGRAPH*, 2003, pp. 587–594.

[3] M. A. G. Ballester, M. G. Linguraru, M. R. Aguirre, and N. Ayache, "On the adequacy of principal factor analysis for the study of shape variability," in *Proc. SPIE Medical Imaging: Image Processing*, J. M. Fitzpatrick and J. M. Reinhardt, Eds., vol. 5747, 2005.

[4] F. L. Bookstein, *Morphometric tools for landmark data: Geometry and Biology*. New York: Cambridge Univ. Press, 1991.

[5] O. T. Carmichael, P. M. Thompson, R. A. Dutton, A. Lu, S. E. Lee, J. Y. Lee, L. H. Kuller, O. L. Lopez, H. J. Aizenstein, C. C. Meltzer, Y. Liu, A. W. Toga, and J. T. Becker, "Mapping ventricular changes related to dementia and mild cognitive impairment in a large community-based cohort," in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2006.

[6] C. Chennubhotla and A. Jepson, "Sparse PCA: Extracting multi-scale structure from data," in *Proc. ICCV*. Vancouver: IEEE, 2001, pp. 641–647.

[7] T. F. Cootes, A. Hill, C. J. Taylor, and J. Haslam, "The use of active shape models for locating structures in medical images," *Image and Vision Computing*, vol. 12, no. 6, pp. 355–366, July 1994.

[8] D. Curran-Everett, "Multiple comparisons: philosophies and illustrations," *Am J Physiol Regul Integr Comp Physiol*, vol. 279, no. 1, pp. R1–R8, July 2000.

[9] S. R. Frost, L. F. Marcus, F. L. Bookstein, D. P. Reddy, and E. Delson, “Cranial allometry, phylogeography, and systematics of large-bodied papionins (primates: Cercopitihecinae) inferred from geometric morphometric analysis of landmark data,” *The Antatomical Record Part A*, vol. 275A, pp. 1048–1072, 2003.

[10] D. Ghosh and N. Amenta, “Landmark transfer using deformable models,” Department of Computer Science, University of California, Davis, Tech. Rep. CSE-2007-6, 2007.

[11] M. Habib, D. Gayraud, A. Oliva, J. Regis, G. Salamon, and R. Khalil, “Effects of handedness and sex on the morphology of the corpus callosum: a study with brain magnetic resonance imaging,” *Brain Cogn.*, vol. 16, no. 1, pp. 41–61, May 1991.

[12] K. L. Narr, P. M. Thompson, T. Sharma, J. Moussai, R. Blanton, B. Anvar, A. Edris, R. Krupp, J. Rayman, M. Khaledy, and A. W. Toga, “Three-dimensional mapping of temporo-limbic regions and the lateral ventricles in schizophrenia: Gender effects,” *Biol. Psychiatry*, vol. 50, pp. 84–97, 2001.

[13] S. Pizer, D. Fritsch, P. Yushkevich, V. Johnson, E. Chaney, and G. Gerig, “Segmentation, registration, and measurement of shape variation via image object shape,” *IEEE Trans. Med. Imaging*, vol. 18, no. 10, pp. 851–865, October 1999.

[14] W. Press, S. Teukolsky, Vetterling, and B. Flannery, *Numerical Recipes In C*, 2nd ed. Cambridge University Press, 1992.

[15] K. Sjöstrand, M. B. Stegmann, and R. Larsen, “Sparse principal component analysis in medical shape modeling,” in *Proc. SPIE Medical Imaging: Image Processing*, J. M. Reinhardt and J. P. W. Pluim, Eds., vol. 6144, 2006.

[16] M. B. Stegmann, K. Sjöstrand, and R. Larsen, “Sparse modeling of landmark and texture variability using the orthomax criterion,” in *Proc. SPIE Medical Imaging: Image Processing*, J. M. Reinhardt and J. P. W. Pluim, Eds., vol. 6144, 2006.

[17] H. Steinmetz, L. Jancke, A. Kleinschmidt, G. Schlaug, J. Volkmann, and Y. Huang, “Sex but no hand difference in the isthmus of the corpus callosum,” *Neurology*, vol. 42, no. 4, pp. 749–752, April 1992.

[18] P. Thompson, R. Dutton, K. Hayashi, A. Lu, S. Lee, J. Lee, O. Lopez, H. Aizenstein, A. Toga, and J. Becker, “3d mapping of ventricular and corpus callosum abnormalities in hiv/aids,” *NeuroImage*, vol. 31, no. 1, pp. 12–23, May 2006.

[19] P. Thompson, K. Hayashi, G. de Zubicaray, A. Janke, S. Rose, J. Semple, M. Hong, D. Herman, D. Gravano, D. Doddrell, and A. Toga, “Mapping hippocampal and ventricular change in Alzheimer’s disease,” *NeuroImage*, vol. 22, no. 4, pp. 1754–1766, August 2004.

[20] M. Üzümcü, A. Frangi, M. Sonka, J. Reiber, and B. Lelieveldt, “ICA vs. PCA active appearance models: Application to cardiac mr segmentation,” in *Proc. MICCAI*, 2003, pp. 451–458.

[21] J. Vermaak and P. Perez, “Constrained subspace modeling,” in *Conf. Computer Vision and Pattern Recog, CVPR’03*, Madison, Wisconsin, June 2003.

[22] S. Witelson, “Hand and sex differences in the isthmus and body of the corpus callosum: a postmortem morphological study,” *Brain*, vol. 112, pp. 799–835, 1989.

[23] H. Zou, T. Hastie, and R. Tibshirani, “Sparse principal component analysis,” *Journal of Computational and Graphical Statistics*, vol. 15, no. 2, pp. 265–286, June 2006.