
Framework for the Statistical Shape Analysis of Brain Structures using SPHARM-PDM

Release 1.00

Martin Styner^{1,2}, Ipek Oguz¹, Shun Xu¹, Christian Brechbühler, Dimitrios Pantazis³, James J Levitt⁴, Martha E Shenton⁴, Guido Gerig^{1,2}

December 6, 2006

¹Dept. Computer Science, Univ. of North Carolina, Chapel Hill NC, USA

²Dept. Psychiatry, Univ. of North Carolina, Chapel Hill NC, USA

³Signal & Image Processing Inst., University of Southern California, Los Angeles, CA

⁴Dept. Psychiatry, VA Boston Healthcare System, Brockton Division, Harvard, Boston MA, USA

Abstract

Shape analysis has become of increasing interest to the neuroimaging community due to its potential to precisely locate morphological changes between healthy and pathological structures. This manuscript presents a comprehensive set of tools for the computation of 3D structural statistical shape analysis. It has been applied in several studies on brain morphometry, but can potentially be employed in other 3D shape problems. Its main limitations is the necessity of spherical topology.

The input of the proposed shape analysis is a set of binary segmentation of a single brain structure, such as the hippocampus or caudate. These segmentations are converted into a corresponding spherical harmonic description (SPHARM), which is then sampled into a triangulated surfaces (SPHARM-PDM). After alignment, differences between groups of surfaces are computed using the Hotelling T^2 two sample metric. Statistical p-values, both raw and corrected for multiple comparisons, result in significance maps. Additional visualization of the group tests are provided via mean difference magnitude and vector maps, as well as maps of the group covariance information.

The correction for multiple comparisons is performed via two separate methods that each have a distinct view of the problem. The first one aims to control the family-wise error rate (FWER) or false-positives via the extrema histogram of non-parametric permutations. The second method controls the false discovery rate and results in a less conservative estimate of the false-negatives.

Contents

1	Introduction	2
2	Methods	3
2.1	SPHARM-PDM	3
	Mathematics behind SPHARM	4
	SPHARM Correspondence	5
2.2	Area-preserving Spherical Mapping	6
2.3	Surface Models from SPHARM: SPHARM-PDM	7

2.4	Alignment and Scaling Normalization	8
2.5	Local Testing of Group Mean Difference using Hotelling T^2 metric	8
2.6	Correction for Multiple Comparison	10
	Minimum over Non-parametric permutation tests	10
	False Discovery Rate	10
3	Computation scheme and Implementation	11
3.1	Processing of Binary Labels: SegPostProcess	12
3.2	Spherical Parametrization: GenParaMesh	12
3.3	SPHARM-PDM: ParaToSPHARMMesh	12
3.4	Statistics: StatNonParamTestPDM	14
3.5	Visualization using KWMeshVisu	16
3.6	Download and Installation	16
4	Results	17
5	Conclusion	20
6	Acknowledgment	20

1 Introduction

Quantitative morphologic assessment of individual brain structures is often based on volumetric measurements. Volume changes are intuitive features as they might explain atrophy or dilation due to illness. On the other hand, structural changes at specific locations are not sufficiently reflected in volume measurements. Shape analysis has thus become of increasing interest to the neuroimaging community due to its potential to precisely locate morphological changes between healthy and pathological structures.

One of the first and most influential research in shape analysis was presented by D’Arcy [1] in his groundbreaking book *On Growth and Form*. In more recent years, several researchers proposed shape analysis via deformable registration to a template ([2, 3, 4, 5]). Inter-subject comparisons are made by analyzing the individual deformable transformations. This analysis of the transformation fields has to cope with the high dimensionality of the transformation, the template selection problem and the sensitivity to the initial position. Nevertheless, several studies have shown stable shape analysis results. [6] and [7] presented some of the first mathematical methods for 3D shape analysis based on sampled descriptions. The shape analysis of densely sampled 3D Point Distribution Models (PDM) and their deformations was first investigated by [8]. Inspired by their experiments, [9] proposed shape analysis based on a parametric boundary description called SPHARM ([10]). The SPHARM shape analysis approach was extended by [11] to use the implied sampled surface (SPHARM-PDM), a method also used by [12]. [13, 14] and Golland [15] proposed shape analysis on medial shape descriptions in 3D and 2D, respectively. They used a fixed topology sampled model with implicit correspondence that is fitted to the objects.

This manuscript presents a comprehensive set of tools that form a pipeline for the computation of 3D structural statistical shape analysis on the object boundary via SPHARM (see Figure 1). It has been applied in several studies on brain morphometry, but can potentially be employed in other 3D shape analysis problems. The main limitations of our shape analysis methodology is the necessity of spherical topology.

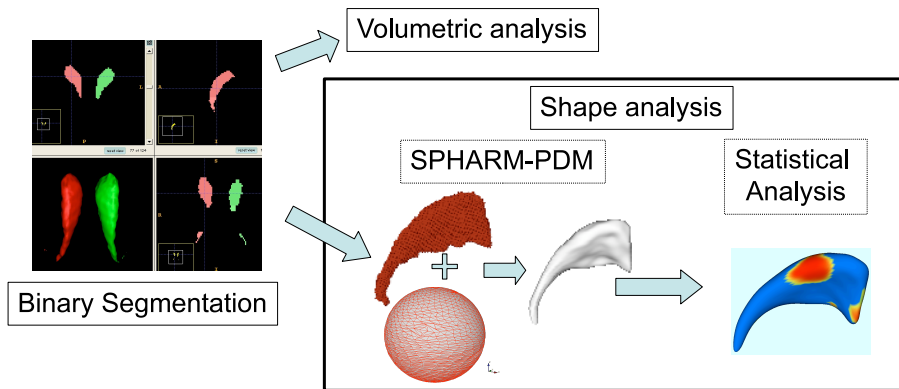


Figure 1: Schematic view of our SPHARM-PDM based shape analysis pipeline.

2 Methods

This section describes the SPHARM-PDM shape description, the statistical analysis methodology, as well as issue of alignment and scaling normalization.

In summary, the input of the proposed shape analysis is a set of binary segmentation of a single brain structure, such as the hippocampus or caudate. These segmentations are first processed to fill any interior holes and a minimal smoothing operation. The processed binary segmentations are converted to surface meshes, and a spherical parametrization is computed for the surface meshes using an area-preserving, distortion minimizing spherical mapping. The SPHARM description is computed from the mesh and its spherical parametrization. Using the first order ellipsoid from the spherical harmonic coefficients, the spherical parametrizations are aligned to establish correspondence across all surfaces. The SPHARM description is then sampled into a triangulated surfaces (SPHARM-PDM) via icosahedron subdivision of the spherical parametrization. These SPHARM-PDM surfaces are all spatially aligned using rigid Procrustes alignment. Group differences between groups of surfaces are computed using the standard robust Hotelling T^2 two sample metric. Statistical p-values, both raw and corrected for multiple comparisons, result in significance maps. Additional visualization of the group tests are provided via mean difference magnitude and vector maps, as well as maps of the group covariance information.

The correction for multiple comparisons is performed via two separate methods that each have a distinct view of the problem. The first one aims to control the family-wise error rate (FWER) or false-positives via the extrema histogram of non-parametric permutations. The second method controls the false discovery rate and results in a less conservative estimate of the false-negatives.

2.1 SPHARM-PDM

In summary, the SPHARM description is a hierarchical, global, multi-scale boundary description that can only represent objects of spherical topology ([10]). The spherical parameterization is computed via optimizing an equal area mapping of the 3D voxel mesh onto the sphere and minimizing angular distortions. The basis functions of the parameterized surface are spherical harmonics. Each individual SPHARM description is composed of a set of coefficients, weighting the basis functions. [16] demonstrated that SPHARM can be used to express shape deformations. Truncating the spherical harmonic series at different degrees results in object representations at different levels of detail. SPHARM is a smooth, accurate fine-scale shape repre-




Figure 2: The SPHARM shape description of a human lateral ventricle shown at 4 different degrees (1, 3, 6, 10 harmonics).

sensation, given a sufficiently high representation level. Based on a uniform icosahedron-subdivision of the spherical parameterization, we obtain a Point Distribution Model (PDM).

Mathematics behind SPHARM

This section discusses the mathematical properties of the spherical harmonic basis functions. It gives a summary of spherical harmonic descriptors as they are presented in Brechbühler's dissertation [17].

Spherical harmonic basis functions Y_l^m , $-l \leq m \leq l$ of degree l and order m are defined on $\theta \in [0; \pi] \times \phi \in [0; 2\pi)$ by the following definitions:

$$Y_l^m(\theta, \phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos \theta) e^{im\phi} \quad (1)$$

$$Y_l^{-m}(\theta, \phi) = (-1)^m Y_l^{m*}(\theta, \phi), \quad (2)$$

where Y_l^{m*} denotes the complex conjugate of Y_l^m and P_l^m the associated Legendre polynomials

$$P_l^m(w) = \frac{(-1)^m}{2^l l!} (1-w^2)^{\frac{m}{2}} \frac{d^{m+l}}{dw^{m+l}} (w^2-1)^l. \quad (3)$$

Table 1 lists explicit expressions for the spherical harmonic functions up to degree 3. The Cartesian notion reveals that the spherical harmonics are polynomials in the 3D space (u_0, u_1, u_2) .

To express a surface using spherical harmonics, the three coordinate functions are decomposed and the surface $\vec{v}(\theta, \phi) = (x(\theta, \phi), y(\theta, \phi), z(\theta, \phi))^T$ takes the form

$$\vec{v}(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l \vec{c}_l^m Y_l^m(\theta, \phi), \quad (4)$$

l	$m = 0$	$m = 1$	$m = 2$	$m = 3$
0	1			
1	$\cos \theta$	$e^{i\phi} \sin \theta$		polar
2	$-1 + 3 \cos^2 \theta$	$e^{i\phi} \cos \theta \sin \theta$	$e^{2i\phi} \sin^2 \theta$	
3	$-3 \cos \theta + 5 \cos^3 \theta$	$e^{i\phi} (1 - 5 \cos^2 \theta) \sin \theta$	$e^{2i\phi} \cos \theta \sin^2 \theta$	$e^{3i\phi} \sin^3 \theta$
0	1			
1	u_2	$u_0 + i u_1$		Cartesian
2	$-1 + 3 u_2^2$	$(u_0 + i u_1) u_2$	$(u_0 + i u_1)^2$	
3	$-3 u_2 + 5 u_2^3$	$(u_0 + i u_1)(1 - 5 u_2^2)$	$(u_0 + i u_1)^2 u_2$	$(u_0 + i u_1)^3$
0	$1/\sqrt{4\pi}$			
1	$\sqrt{3/4\pi}$	$\sqrt{3/8\pi}$		
2	$\sqrt{5/16\pi}$	$\sqrt{15/8\pi}$	$\sqrt{15/32\pi}$	
3	$\sqrt{7/16\pi}$	$\sqrt{21/64\pi}$	$\sqrt{105/32\pi}$	$\sqrt{35/64\pi}$

Table 1: Explicit expressions of the spherical harmonics up to degree 3, in both polar and Cartesian form due to Brechbühler. The last part of the table gives the common normalizing constants, e.g. $Y_1^0 = \sqrt{3/4\pi}u_2$.

where the coefficients \vec{c}_l^m are three-dimensional vectors due to the three coordinate functions. The coefficients \vec{c}_l^m are obtained by solving a least-squares problem. Therefore, the values of the basis functions are gathered in the matrix $Z = (z_{i,j(l,m)})$ with $z_{i,j(l,m)} = Y_l^m(\theta_i, \phi_i)$, where $j(l,m)$ is a function assigning an index to every pair (l,m) and i denotes the indices of the n_{vert} points to be approximated. The coordinates of these points are arranged in $V = (\vec{v}_1, \vec{v}_2, \dots, \vec{v}_{n_{vert}})^T$ and all coefficients are gathered in $C = (\vec{c}_0^0, \vec{c}_1^{-1}, \vec{c}_1^0, \dots)^T$. The coefficients that best approximate the points in least-squares sense are obtained by

$$C = (Z^T Z)^{-1} Z^T V. \quad (5)$$

Using spherical harmonic basis functions, we obtain a hierarchical surface description that includes further details as more coefficients are considered. This is illustrated in Fig. 2.

SPHARM Correspondence

Correspondence of SPHARM-PDM is determined by normalizing the alignment of the spherical parameterization to an object-specific frame. In the studies presented in this paper, the normalization is achieved by rotation of the parameterization, such that the spherical equator, 0° and 90° longitudes coincide with those of the first order ellipsoid (see Figure 3). After this parametrization normalization, corresponding surface points across different objects possess the same parameterization.

The SPHARM-PDM shape analysis is visualized in Figure 1 using a caudate structure. Prior to the shape analysis, the group average object is computed for each subject group, and an overall average object is computed over all group average objects. Each average structure is computed by averaging the 3D coordinates of corresponding surface points across the group. The overall average object is then used in the shape analysis as the template object. At every boundary point for each object, we compute a distance map representing the signed local Euclidean surface distance to the template object. The sign of the local distance is computed using the direction of the template surface normal. In the global shape analysis, the average of the local distances across the whole surface is analyzed with a standard group mean difference test. The local shape analysis is computed by testing the local distances at every boundary point. This results in a significance map that represents the significance of these local statistical tests and thus allows locating significant shape differences between the groups. We corrected the shape analysis for the multiple comparison problem using a uniformly sensitive, non-parametric permutation test approach ([18]). The non-corrected significance map




Figure 3: Left: Illustration of SPHARM correspondence via alignment of the spherical parametrization using the first order ellipsoid meridian and equator. Right: Set of 6 caudate structures with correspondence shown at selected locations via colored spheres.

is an optimistic estimate of the real significance, whereas the corrected significance map is a pessimistic estimate that is guaranteed to control the rate of false positives at the given level α (commonly $\alpha = 0.05$) across the whole surface.

We have compared the SPHARM correspondence to other types of correspondences and it compared favourably to human expert established correspondence [19].

2.2 Area-preserving Spherical Mapping

The appropriate parameterization of the points of a surface description is a key problem for correspondence finding, as well as for an efficient SPHARM representation. Every point i on the surface is to be assigned a parameter vector (θ_i, ϕ_i) that are located on the unit sphere. A homogeneous distribution of the parameter space is essential for an efficient decomposition of the surface into SPHARM coefficients. We give here a brief summary of the surface parameterization procedure proposed by Brechbühler [10].

A bijective mapping of the surface to the unit sphere is created, i. e., every point on the surface is mapped to exactly one point on the sphere, and vice versa. The main idea of the procedure is to start with an initial parameterization. This initial parameterization is optimized so that every surface patch gets assigned an area in parameter space that is proportional to its area in object space. The initial mapping from surface to parameter space is constructed using discrete Laplace's equations to solve the corresponding Dirichlet problem. To obtain a homogeneous distribution of the parameter space over the surface, this initial parameterization is modified in a constrained optimization procedure considering two criteria:

1. Area preservation: Every object region must map to a region of proportional area in parameter space.
2. Minimal distortion: Every quadrilateral should map to a spherical quadrilateral in parameter space.

Recursive Level	Linear Level	Number of Surface Points
1	1	12
2	2	42
3	4	162
4	8	642
5	16	2562

Table 2: Comparison of subdivision levels for recursive and linear subdivision of the icosahedron.

Brechtbühler establishes constraints for area preservation, while the distortion of the mesh serves as the objective function during optimization. The optimization solves the resulting system of nonlinear equations by linearizing them and taking Newton steps.

2.3 Surface Models from SPHARM: SPHARM-PDM

From the SPHARM description we can compute triangulated surfaces by sampling the spherical parameterization uniformly on the sphere. Equidistant sampling in the parameter space leads to a dense sampling around the poles ($\theta = 0, \theta = \pi$) and a coarse sampling around the equator ($\theta = \pi/2$). This fact can be explained by the poles being mapped to all points having $\phi = 0 \dots 2\pi$ and $\theta = 0$, or $\theta = \pi$ (see also Fig. 3). Using a linear, uniform icosahedron subdivision shown in Fig. 4, however, we gain a good approximation of a homogeneous sampling of the spherical parameter space and thus also of the object space. The subdivision is linear in the number of subdivisions along each edge of the original icosahedron, rather than the commonly used recursive subdivisions. Any recursive subdivision has a corresponding linear subdivision, but most linear subdivisions have no corresponding recursive subdivision (see Table 2).




Figure 4: Icosahedron subdivision for different levels of sampling. From left to right: Base icosahedron, subdivision factors 2, 4 and 6.

Using the pre-computed parameter locations (θ_i, ϕ_i) from the icosahedron subdivision, we can compute the PDM directly from the coefficients, as the parameter locations stay constant for all objects at a given subdivision level. The sampling points \vec{x}_i at the locations (θ_i, ϕ_i) are obtained using equation (4):

$$\vec{x}_i = \sum_{l=0}^K \sum_{m=-l}^l \vec{c}_l^m Y_l^m(\theta_i, \phi_i).$$

In our experience the sampling for the surface of brain structures should be chosen between subdivision level 10 (1002 surface points) and 20 (4002 surface points) depending on the complexity of the objects. For

example, hippocampal surfaces are well represented by a subdivision level 10 and caudate surfaces by a subdivision level 15.

2.4 Alignment and Scaling Normalization

Alignment and scaling of the objects are two important issues in shape analysis that are also discussed here [11]. In general, we usually chose rigid-body Procrustes alignment ([20]) followed by no scaling normalization or scaling inversely to the intra-cranial cavity volume (ICV). Commonly both results are reported in a shape analysis study, the original scale analysis represents an unbiased raw analysis, whereas the ICV scale analysis corrects for differences in overall interior head size and thus approximatively normalizes for gender and age differences. In general both types of analysis should yield similar patterns of differences and special care needs to be taken if this is not the case.




Figure 5: Left: Schematic view of testing procedure: A set of features per surface point (shown as colormaps) are analyzed using uni- or multi-variate statistics, resulting in a raw significance map (bottom) .Right: The raw significance maps are overly optimistic and need to be corrected to control for the multiple comparison problem. The corrected significance maps on the other hand are commonly pessimistic estimates.

2.5 Local Testing of Group Mean Difference using Hotelling T^2 metric

After shape description, correspondence establishment, alignment and scaling normalization, the next step in the shape analysis pipeline consists of testing for differences between groups at every surface location (see Figure 5). This can be done in 2 main fashions:

- *Analyzing the magnitude of the local difference vector to a template:* For this option, a template needs to be first selected, usually this is the common mean of the 2 groups or the mean of separate control group. The magnitude of the difference is easily understood and results in difference maps for each

subject. Also, the resulting univariate statistical analysis is quite well known and local significance can be easily computed using the Student t metric. The main disadvantage of this method is the need to select a template, which introduces an additional bias into the statistical analysis. We applied this method in earlier studies of hippocampi [21] and ventricles [22].

- *Analyzing the spatial location of each point:* For this option, no template is necessary and multivariate statistics of the (x, y, z) location is necessary. We have chosen to use the Hotelling T^2 two sample difference metric as a measurement of how 2 groups locally differ from each other. The standard Hotelling T^2 is defined as $T^2 = (\mu_1 - \mu_2)'(\Sigma(\frac{1}{n_1} + \frac{1}{n_2}))^{-1}(\mu_1 - \mu_2)$, where $\Sigma = (\Sigma_1(n_1 - 1) + \Sigma_2(n_2 - 1))/(n_1 + n_2 - 2)$ is the pooled covariance matrix. An alternative modified Hotelling T^2 metric is less sensitive to group differences of the covariance matrixes and the number of samples[23]: $T^2 = (\mu_1 - \mu_2)'(\Sigma_1\frac{1}{n_1} + \Sigma_2\frac{1}{n_2})^{-1}(\mu_1 - \mu_2)$. All our current studies are based on this modified Hotelling T^2 metric.

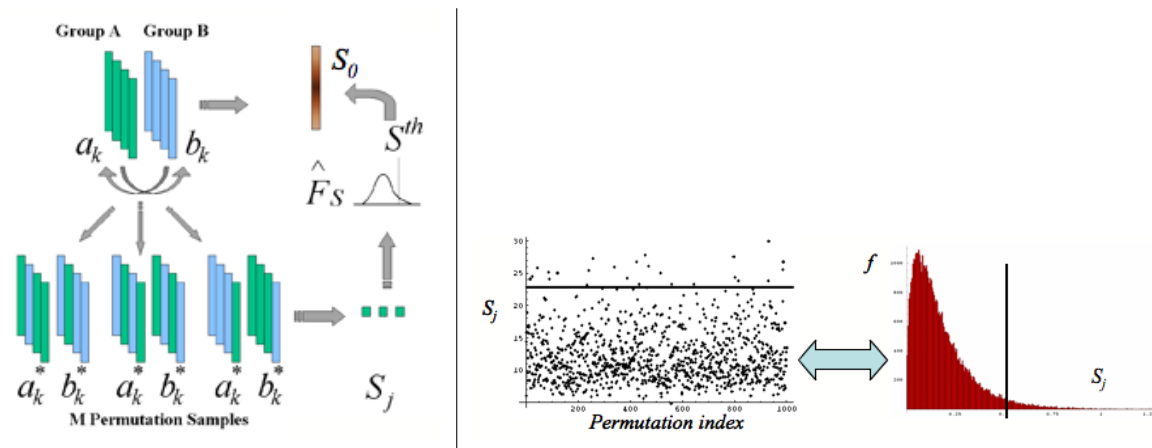


Figure 6: Left: Scheme of P-value computation via permutation tests. The real group difference S_0 (e.g. Hotelling T^2) is compared to that the group differences S_j of random permutations of the group labels. The quantile in the S_j histogram associated with S_0 is the p-value. Right: Example run on test data with 1000 random permutations. On the left the individual S_j are plotted and on the left the corresponding S_j histogram is shown. The black line indicates the value of S_0 .

We then want to test the two groups for differences in the means of the selected difference metric (univariate: Student t, multivariate: Hotelling T^2) at each spatial location. Permutation tests are a valid and tractable approach for such an application, as they rely on minimal assumptions and can be applied even when the assumptions of the parametric approach are untenable. Non-parametric permutation tests are exact, distribution free and adaptive to underlying correlation patterns in the data. Further, they are conceptually straightforward and, with recent improvements in computing power, are computationally tractable.

Our null hypothesis is that the distribution of the locations at each spatial element is the same for every subject regardless of the group. Permutations among the two groups satisfy the exchangeability condition, i.e. they leave the distribution of the statistic of interest unaltered under the null hypothesis. Given n_1 members of the first group $a_k, k = 1 \dots n_1$ and n_2 members of the second group $b_k, k = 1 \dots n_2$, we can create $M \leq \binom{n_1 + n_2}{n_2}$ permutation samples. A value of M from 20000 and up should yield results that are negligibly different from using all permutations [24].

2.6 Correction for Multiple Comparison

The local shape analysis involves testing from a few to many thousands of hypothesis (one per surface element) for statistically significant effects. It is thus important to control for the multiple testing problem (see Figure 5), and the most common measure of multiple false positives is the familywise error rate (FWER). The multiple testing problem has been an active area of research in the functional neuroimaging community. The most widely used methods in the analysis of neuroimaging data use random field theory [25] [26] and make inferences based on the maximum distribution. In this framework, a closed form approximation for the tail of the maximum distribution is available, based on the expected value of the Euler characteristic of the thresholded image [26]. However, random field theory relies on many assumptions such as the same parametric distribution at each spatial location, a point spread function with two derivatives at the origin, sufficient smoothness to justify the application of the continuous random field theory, and a sufficiently high threshold for the asymptotic results to be accurate.

In our work, we are employing non-parametric permutation tests and false discovery rate as two alternative correction methods for the multiple comparison problem.

Minimum over Non-parametric permutation tests

Non-parametric methods deal with the multiple comparisons problem [18] and can be applied when the assumptions of the parametric approach are untenable. Similar methods have been applied in a wide range of functional imaging applications [27, 28, 29].

Details of this method are described in [18] and only a summary is provided here. The correction method is based on computing first the local p-values using permutation tests. The minimum of these p-values across the surface is then computed for every permutation. The appropriate corrected p-value at level α can then be obtained by the computing the value at the α -quantile in the histogram of these minimum values. Using the minimum statistic of the p-values, this method correctly controls for the FWER, or the false positives, but no control of the false negatives is provided. The resulting corrected local significance values can thus be regarded as pessimistic estimates akin to a simple Bonferroni correction.

False Discovery Rate

Additionally to the non-parametric permutation correction, we have also implemented and applied a False Discovery Rate Estimation (FDR) [30, 31] method. The innovation of this procedure is that it controls the expected proportion of false positives only among those tests for which a local significance has been detected. The FDR method thus allows an expected proportion (usually 5%) of the FDR corrected significance values to be falsely positive. The correction using FDR provides an interpretable and adaptive criterion with higher power than the non-parametric permutation tests. FDR is further simple to implement and computationally efficient even for large datasets.

The FDR correction is computed as follows [31]:

1. Select the desired FDR bound q , e.g. 5%. This is the maximum proportion of false positives among the significant tests that you are willing to tolerate (on average).
2. Sort the p-values smallest to largest.
3. Let p_q be the p-value for the largest index i of the sorted p-values $p_{\text{sort},i} \leq q \cdot i/N$, where N is the number of vertices.

4. Declare all locations with a p-value $p \leq p_q$ significant.

3 Computation scheme and Implementation

This section describes the set of shape analysis tools and the use of these in succession to compute the whole shape analysis (see Figure 7). The general workflow starts from the segmented datasets with binary labels for each structure of interest. The *SegPostProcess* program extracts a single binary label and applies heuristic methods to ensure the spherical topology of the segmentation. In the next step, the *GenParaMesh* computes the surface mesh corresponding to the segmentation, as well as a spherical parametrization of this surface mesh via area preserving mapping (see Section 2.2). The SPHARM-PDM shape representation is then computed using the *ParaToSPHARMMesh* program, which also handles issues of alignment. These three programs have to be run on all datasets in a shape analysis study. Often they are applied in parallel to the datasets. After all SPHARM-PDM representations have been computed, the statistical testing program *StatNonParamTestPDM* computes the difference maps, the significance maps, as well as the descriptive statistics such as the mean surfaces and covariance matrixes,

All these tools are command line tools without any graphical user interface. They have an online help/usage describing the input and output parameter specification. This information is printed on the terminal window by calling the program without any parameters or using the '-help' option.

All input and output meshes are stored in ITK-readable format of an `itk::MeshSpatialObject`. Conversion tools to/from other formats, such as OpenInventor and STL are also provided.

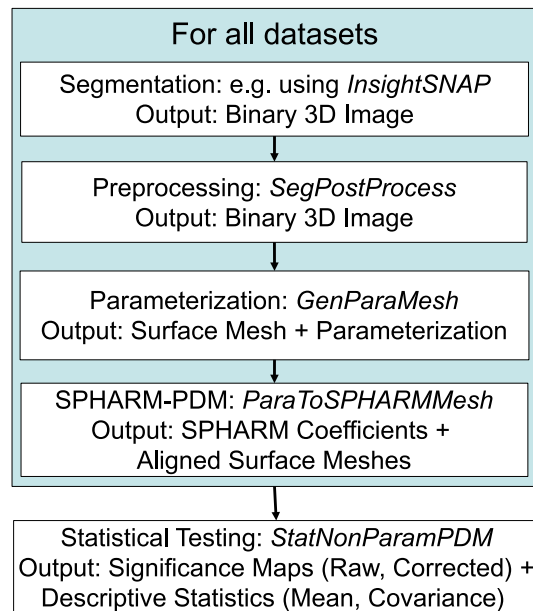


Figure 7: Workflow of the shape analysis pipeline starting at the segmentation of the brain structures to the local final statistical analysis.

3.1 Processing of Binary Labels: SegPostProcess

The application of the *SegPostProcess* program is optional and but we recommend its use, as it ensure spherical topology of the segmentation and offers a series of services:

1. Extraction of a single label or a label range: This is especially helpful if the label dataset contains multiple labels.
2. Resampling of the label data: Segmentation data from MRI or other modalities are often stored at rather coarse resolution or at non-isotropic resolution. The input to *GenParaMesh* has to be of isotropic resolution and a relative fine resolution is recommended. In our studies, we have commonly resampled the data to an isotropic resolution of 0.5 mm to 0.75 mm.
3. Spherical topology: Any cavities fully enclosed in the segmentation is filled, two smoothing operations are applied, a full 6-connectedness is ensured, and only the largest component in the image is extracted. The cavity filling only fills fully interior structures and thus leaves the exterior boundary unedited. The first smoothing operation is implemented via a binary closing operation and the second smoothing operation is a level set based anti-alias smoothing (ITK filter `itk::AntiAliasBinaryImageFilter` [32]).

A common use of this tool is illustrated by the following command line. The output is a single label dataset.

```
SegPostProcess Label.hdr -o Label_PP.hdr -label 1-space 0.75,0.75,0.75
```

3.2 Spherical Parametrization: GenParaMesh

GenParaMesh first extracts the surface of the input label segmentation by following the 'cracks' between the voxels of the foreground (label) and the background. The resulting mesh coordinates are thus off by minus half a voxel-width from coordinate systems that place the origin at the center of the voxels. This surface mesh is mapped to sphere using the methods described in Section 2.2. If the tool reports a bad Euler number, then this shows that the extracted surface is not of spherical topology, whereas the vice versa is not necessarily true. An indication of a successful parametrization is given by the output line 'initially active:'. If the number of initially active equations is low, a successful computation is likely. In our experience, less than 2% of the datasets do not finish successfully when *SegPostProcess* is first applied with a reasonably high resolution.

A common use of this tool is illustrated by the following command line. The output is two meshes, one for the surface and one for the spherical parametrization.

```
GenParaMesh Label_PP.hdr -iter 800 -label 1
```

3.3 SPHARM-PDM: ParaToSPHARMMesh

This tool computes the SPHARM-PDM representation and resolves issues of correspondence and alignment. The input is a surface mesh with spherical parametrization. It is not necessary that this parametrization has been computed using our *GenParaMesh*, but other parametrizations can equally be used, such as conformally mapped spherical parametrizations. As a first step, the raw spherical harmonic coefficients are computed only up to the first degree and the first order ellipsoid is determined. The spherical parametrization is then

rotated such that the poles of the first order ellipsoid are coinciding with the poles of the parametrization. The spherical harmonic coefficients are recomputed up to the degree specified on the command line option. Additionally, the surface points of the SPHARM are reconstructed using the icosahedron subdivision scheme described in section 2.3.

The two main parameters of this tool are the maximal degree for the SPHARM computation and the subdivision level for the icosahedron subdivision. In our experience, the SPHARM maximal degree is should be chosen between 12 (hippocampus) to 15 (lateral ventricle, caudate) for brain structures. If the degree is chosen too high the reconstructed SPHARM surface will often show signs of voxelization. By using the '-paraOut' command line option, the spherical icosahedron subdivision, as well as local phi and theta attribute files for the quality control visualization with *KWMeshVisu* are written out.




Figure 8: Quality Control visualization of the SPHARM correspondence with *KWMeshVisu* using the Phi-attribute file shown on 3 randomly selected cases. Same color represent the same ϕ parameter value of the spherical parametrization.

Correspondence: The SPHARM shape description has an object-inherent correspondence based on the first order ellipsoid. Unfortunately this is not fully unique as the first order ellipsoid can be flipped along any of its axis with the same result. Thus, without additional measures the correspondence is ambiguously defined in regard to flips along these axis. In order to clarify this ambiguity, an additional SPHARM object (described by SPHARM coefficients) can be provided as a flip-template. This flipTemplate is used to test all possible flips of the parametrization along the first order ellipsoid axis and select the one whose reconstruction has minimal distance to the flip-template. Commonly, the first object is computed without providing a flip-template and then serves for all subsequent objects as the flip-template (see also Section 4).

Alignment: The *ParaToSPHARMMesh* tool provides two types of alignment. The first type performs an alignment of the normalized first order ellipsoid axis to the unit-axis with the ellipsoid center located at the origin. The second alignment type is a rigid-body Procrustes alignment [6], i.e. Procrustes alignment only with translation and rotation. Whereas the first alignment is object inherent, the Procrustes alignment is in need of a template, which can optionally be supplied. It is quite important in case of Procrustes alignment that a corresponding flip-template is supplied.

A common use of this tool is illustrated by the following command line. The output is a series of SPHARM coefficients and SPHARM-PDM meshes, one set in the original coordinate system, one in the first order ellipsoid aligned coordinate system and one in the Procrustes aligned coordinate system.

```
ParaToSPHARMMesh Label_surf.meta Label_para.meta -subdivLevel 10 -spharmDegree 12 \
    -flipTemplate template.coef -regTemplate template.meta
```

TroubleShooting: In rare cases, the parametrization alignment using the first order ellipsoid correspondence can result in improperly flipped parametrizations and consequently flipped coordinates in the ellipsoid aligned surfaces. This flipping can occur along any of the axes of the first order ellipsoid. Such flips would

be visible in the QC visualization (a correct, normal QC is shown in Figure 10). In such cases, an optional, user-initiated flip of the parametrization can be performed using the 'FinalFlip' option. The result of a parametrization flip is shown in Figure 9 where the parametrization is flipped along the z-axis and so are the coordinates of the ellipse aligned datasets. The smallest axis is the x-axis, the second smallest is the y-axis, and the long axis of the first ellipsoid is the z-axis,




Figure 9: Single user initiated parametrization flip at the x-axis. Left: original surface with phi coloring, right: flipped surface with phi coloring after ellipsoid alignment. It is clearly visible that flipping the parametrization will lead also to a coordinate flip for the first order ellipsoid aligned surface

3.4 Statistics: StatNonParamTestPDM

Once all SPHARM-PDM meshes have been computed, group differences of the local surface point distributions can be assessed using the *StatNonParamTestPDM*. There are two main types of results from the statistical testing step: a) descriptive group statistics, such as the mean and covariance information, and b) group mean difference hypothesis testing. The descriptive statistics are used for quality control, sanity check, as well as these are necessary to make sense of the computed significance maps ("Is that significant region enlarged in the patient population?"). The significance maps show the regions of significant difference of the Hotelling T^2 metric with 3 different maps (see also section 2.5): the raw p-value map, the FDR corrected p-value map, and the permutation based corrected p-value map. The value of the Hotelling T^2 metric is also useful to visualize the full range of the effect size.

The current version of the tool also computes a single global shape difference by averaging the local Hotelling T^2 metric across the whole surface. The p-value is computed the same way as for the local tests using non-parametric permutation tests. No correction is needed for the global shape difference as only a single test is performed. Additional other measures of global shape difference will be implemented, such as percentile measures of the Hotelling T^2 representing stable estimations of extremal difference, such as the 68% and 95% percentiles.

The main parameter of the testing procedure is the number of permutations employed for both the computation of the raw and corrected p-values. As a rule of thumb, we are employing 10 times the number of surface mesh points, which is a safe estimate and results in quite extensive computation and memory requirements. Other relevant parameters include the maximal significance threshold (commonly chosen at 5%), as well as the resolution of the significance correction captured by the number of thresholds uniformly spaced between null and the maximal significance threshold (a good value is about 1000). For example, a maximal threshold of 5% with 10 threshold steps results in a correction resolution of $5\%/10 = 0.5\%$. The number of steps parameter does not affect the raw significance map, which is rather determined by the number of permutations.

The *StatNonParamTestPDM* tool has several options for specifying its input. There are 2 main input types: a) text file listing all mesh files with corresponding information b) text file containing raw feature information. Using the second option, the tool can be applied to many different types of data that are embedded in a

linear space, but limited descriptive group statistics are provided. The first option is simpler to use and has a richer set of output. Both options use ASCII text files, which reserve a single line per subject. For the mesh-list file, each line contains first the group identifier number (often either 0 or 1), then the scaling factor for an optional scaling normalization, followed by the absolute or relative path to the mesh-file. The scaling normalization is only performed when the appropriate '-scale' command line option is also present, otherwise the scaling information is simply omitted.

The scale information can also be set to 1.0 everywhere if no scaling is desired. In order to normalize for cranial size differences, we often apply a scaling normalization to the internal cranial cavity volume (ICC, usually computed by a separate full brain tissue classification $ICC = WM + GM + CSF$ volumes). The scaling is applied as a divisor. For a simple normalization with ICC, the scaling information would be $scale_i = ICC_i^{1/3}$. Since shape differences are more easily interpreted in the original unit space (usually millimeter units) the scale information is divided by the mean over all ICC values: $scale_i = ICC_i^{1/3} / \mu_{ICC}^{1/3}$. This is equivalent of scaling each structure as if its overall brain is normalized to a standard head size.

In case, data other than meshes are to be tested, the current version of the program allows the testing of uni-variate, bi and tri-variate data via the '-2DvectorData' and '-3DvectorData' options. Additional information needs to be provided, identifying the column numbers (indexed at 0) for the group identification number and for the first feature, as well as the numbers of features.

Output files:

- MeanA.txt, MeanB.txt, MeanAll.txt: The mean surface of the groups A and B (the program prints on the terminal output which label is group and which is group B0, as well as the correctly weighted joint mean surface).
- MeanDiff.txt: The absolute value of the difference between the mean surfaces of the 2 groups. This is the lengths of the vectors in the difference vector field (meanDiffVec.txt)
- meanDiffVec.txt: The difference vector field between the mean surfaces of the 2 groups.
- CovarA.txt, CovarB.txt, CovarAll.txt: The covariance ellipsoid information of the groups, as well as over all subjects. The information is stored as the 3 orthogonal axes of the ellipsoid (largest to smallest), which are equivalent to the product of the square root of the i-th Eigenvalue and the i-th Eigenvector $\sqrt{\lambda_i} * ev_i$ of the covariance metric. The Eigendecomposition of the covariance metric represent the directions of major variance.
- realStat.txt: This stores the Hotelling T^2 mean difference metric
- Tstat.txt: This is the square root of the Hotelling T^2 mean difference metric
- r2stat.txt: This is the local r^2 statistics.
- avgPval.txt: This file contains a single value, the p-value for computing the difference between the group when using the average of the mean surface difference across the whole surface as a single univariate test. This represents the global shape difference.
- rawPval.txt: These are the local uncorrected, raw p-values, which are too optimistic.
- FDRPval.txt: Correction of the p-values using FDR.
- corrPval.txt: Correction of the p-values using the permutation-based correction.

A common use of this tool is illustrated by the following command line. The output is are mean surface meshes for both groups and for all datasets, as well as attribute files for Covariance statistics, mean difference, local significance, and Hotelling T^2 values. Additionally the significance of the global shape difference is written out.

```
StatNonParamTestPDM listfile.txt -out stats -surfList -numPerms 20000 \
  -signLevel 0.05 -signSteps 1000
```

3.5 Visualization using KWMeshVisu

We have also developed a tool for the intuitive visualization of the statistical results, as well as for inspecting the established SPHARM correspondence. This tool is called KWMeshVisu and the SPHARM-PDM and statistical visualizations in this document have been created with it. The use and implementation details of this tool have been described in a separate paper, which has been published at the ISC/NAMIC Workshop on open science at MICCAI 2006.




Figure 10: Quality Control visualization of the SPHARM correspondence with *KWMeshVisu* using the Phi-attribute (left) and Theta-attribute (right) file on the hippocampus sample data downloadable at <http://www.ia.unc.edu/dev/download/shapeAnalysis/>.

3.6 Download and Installation

The shape analysis framework described in this paper is both accessible through the Sandbox of the National Alliance for Medical Image Computing NAMIC or directly as part of the UNC NeuroLib library [33]. Instructions for the download and installation process of the UNC NeuroLib can be found in the tutorial section. For compilation ITK libraries as well as lapack and blas libraries are necessary. Lapack and blas are often installed standardly on UNIX system, or can be downloaded from www.netlib.org/lapack.

- NAMIC: www.na-mic.org/Wiki/index.php/Engineering:SandBox

- UNC: www.ia.unc.edu/dev

Binary compilations for 32 and 64 bit redhat linux distributions can be downloaded from the following URL <http://www.ia.unc.edu/dev/download/shapeAnalysis>. This distribution also contains sample data, and sample scripts.

4 Results

As exemplary test study we applied our shape analysis framework to a schizo-typal personality disorder (SPD) study on the caudate brain structure that was available to us through the NAMIC network [34]. The data consists of structural SPGR MRI datasets from 15 right handed medication-free male SPD patient, as well as from 14 matched healthy control subjects. Although both left and right hemispheric caudate were outlined on the MRI data, only the right caudate is presented here as a test study. Additionally, we are currently developing appropriate 3D synthetic test datasets, which will be published online as open access data.

The following simple script was used for the computation of the SPHARM-PDM descriptions:

```
#!/bin/tcsh -f

foreach case (/Data/caud*/*_R.gipl)
  set ppcase = $case:r_pp.gipl
  SegPostProcess $case -o $ppcase-label 1 -space 0.5,0.5,0.5

  GenParaMesh $ppcase -iter 1000 -label 1

  set meshPrefix = $ppcase:r
  ParaToSPHARMMesh ${meshPrefix}_surf.meta ${meshPrefix}_para.meta \
    -subdivLevel 15 -spharmDegree 15 \
    -flipTemplate flipTemplate.coef -regTemplate regTemplate.meta
end
```

The correspondence of the SPHARM-PDM description was visually controlled (see Figure 8) for all datasets. For the group mean difference testing step, a text files was created listing the names of all mesh-files preceded by the group identifier (either 0 or 1) and the scaling factor. In this case the scaling factor was 1.0 for all cases, i.e. no scaling was performed. A sample line in this mesh list text-file thus looks like the following:

```
1 1.0 /Data/caud1/caud1_R_pp_surfSPHARM_procalign.meta
```

This mesh list file was then the input of the statistical testing program, called in the following fashion:

```
StatNonParamTestPDM caudList.txt -out stats -surfList -numPerms 20000 \
  -signLevel 0.05 -signSteps 1000
```

The descriptive group statistics (see Figure 11) show that the mean objects of the 2 groups are quite similar. The differences between the mean are overall small (max at 1.6 millimeter), but most pronounced in the anterior head region, the superior body region and the posterior tail region ¹. The visualization of the covariance ellipsoids show that the variability is reduced in the body section compared to the anterior head and especially the posterior tail section.

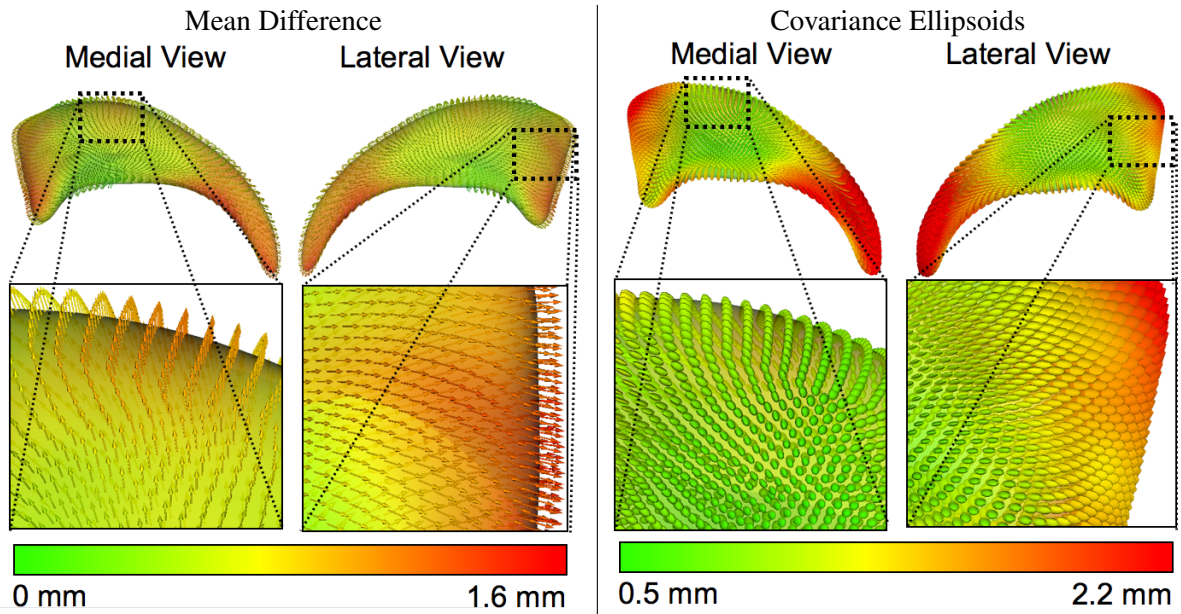


Figure 11: Descriptive statistics of right caudate analysis. The left visualization shows the mean difference as a colormap from green (0 mm difference) to red (1.6 mm difference) as well as a vectorfield on the overall mean surface mesh. For two regions with large differences a zoomed view is provided. The right visualization shows the covariance ellipsoid colored with the magnitude of the largest ellipsoid axis (axis of largest covariance). The same two zoomed views are provided. The variance in the first zoom section is markedly reduced compared to the second zoom section.

The results of the mean difference hypothesis tests are visualized in Figure 12. The raw significance maps display an overly optimistic estimate with the superior body, as well as the anterior head region to be significantly different. The conservative, permutation-based correction map provides a pessimistic estimate with only a single significant region in the anterior head region. The result of the false discovery rate correction is in-between the raw and the permutation-based correction. On one hand, significant parts in the anterior tail section are visible in the raw significant map, but are not preserved in the FDR-corrected map. On the other hand, significant parts in the superior body as well as the anterior head region are present in the FDR-corrected maps. The global shape difference has a p-value of $p = 0.009$.

Based on these results, a significant shape difference between the two populations on the right caudate is clearly suggested. The main differences are enlarged sections in controls compared to SPD located in the anterior head and superior body regions. For a full shape analysis, we recommend to both analyze the objects in their original scale (scaling factor = 1.0, as in this study), as well as normalized for the individual intra-cranial cavity volume (ICV). The ICV is usually acquired in a separate brain tissue segmentation and the scaling factor f is then chosen as $f_i = (Mean(ICV)/ICV_i)^{\frac{1}{3}}$. As this study serves as an example of how to employ our shape analysis framework, the analysis of the ICV-normalized data is omitted here.

¹The naming of caudate head and tail regions in this section are not based on anatomical labels

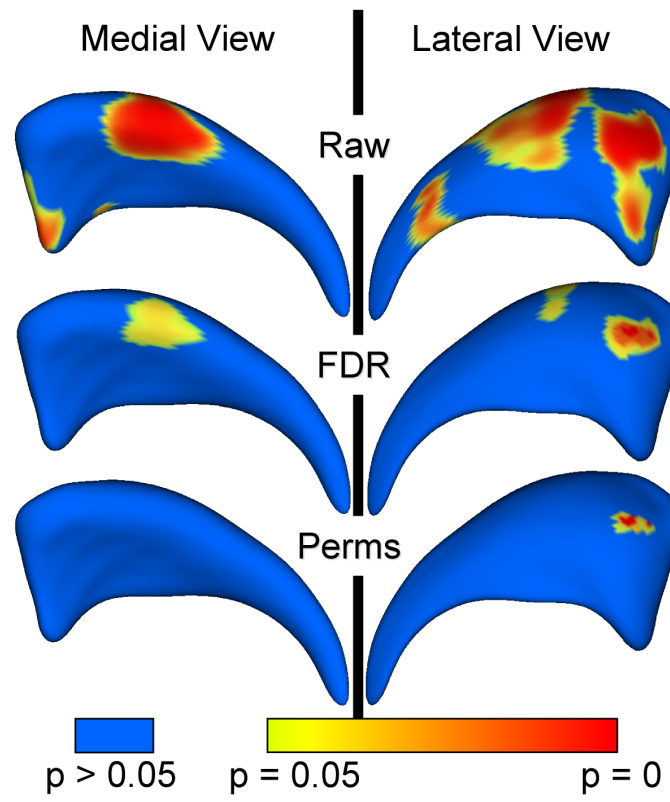


Figure 12: Significance testing on right caudate study

5 Conclusion

We have presented a comprehensive, local shape analysis framework founded on the SPHARM-PDM shape representation and local permutation tests. The whole software is available as open source and is currently in use by several research groups outside of UNC.

We presented research and software development in progress, and many additional enhancements of the current tool base are being developed or are planned. These enhancements include the incorporation of subject covariates, such as age, gender, handedness etc, as well as an user-interface based control tool for the whole shape analysis.

6 Acknowledgment

We are thankful to Steve Pizer and Sarang Joshi for highly valuable discussions in regard to shape analysis. This research has been/is supported by the UNC Neurodevelopmental Disorders Research Center HD 03110, NIH Roadmap for Medical Research, Grant U54 EB005149-01, NIH NIBIB grant P01 EB002779, EC-funded BIOMORPH project 95-0845, Department of Veteran Affairs Merit Award, VA Research Enhancement Award Program, NIH R01 MH50747, and K05 MH070047.

References

- [1] D. Thomson, *On Growth and Form*, Cambridge University Press, second edition, 1942. 1
- [2] C. Davatzikos, M. Vaillant, S. Resnick, J.L Prince, S. Letovsky, and R.N. Bryan, "A computerized method for morphological analysis of the corpus callosum," *J. of Comp. Assisted Tomography*, vol. 20, pp. 88–97, Jan./Feb 1996. 1
- [3] S. Joshi, M Miller, and U. Grenander, "On the geometry and shape of brain sub-manifolds," *Pat. Rec. Art. Intel.*, vol. 11, pp. 1317–1343, 1997. 1
- [4] J.G. Csernansky, S. Joshi, L.E. Wang, J. Haller, M. Gado, J.P. Miller, U. Grenander, and M.I. Miller, "Hippocampal morphometry in schizophrenia via high dimensional brain mapping," *Proc. Natl. Acad. Sci. USA*, vol. 95, pp. 11406–11411, September 1998. 1
- [5] J.G. Csernansky, L. Wang, D. J. Jones, D. Rastogi-Cru, G. Posener, J.A. and Heydebrand, J.P. Miller, U. Grenander, and M.I. Miller, "Hippocampal deformities in schizophrenia characterized by high dimensional brain mapping," *Am. J. Psychiatry*, vol. 159, pp. 1–7, December 2002. 1
- [6] F.L. Bookstein, "Shape and the Information in Medical Images: A Decade of the Morphometric Synthesis," *Comp. Vision and Image Under.*, vol. 66, no. 2, pp. 97–118, May 1997. 1, 3.3
- [7] I.L. Dryden and K.V. Mardia, "Multivariate shape analysis," *Sankhya*, vol. 55, pp. 460–480, 1993. 1
- [8] T. Cootes, C.J. Taylor, D.H. Cooper, and J. Graham, "Active shape models - their training and application," *Comp. Vis. Image Under.*, vol. 61, pp. 38–59, 1995. 1
- [9] G. Gerig, M. Styner, M.E. Shenton, and J.A. Lieberman, "Shape versus size: Improved understanding of the morphology of brain structures," in *MICCAI*, 2001, pp. 24–32. 1

- [10] C. Brechbühler, G. Gerig, and O. Kübler, “Parametrization of closed surfaces for 3-D shape description,” *Comp. Vision, Graphics, and Image Proc.*, vol. 61, pp. 154–170, 1995. [1](#), [2.1](#), [2.2](#)
- [11] G. Gerig, M. Styner, D. Jones, D. Weinberger, and J. Lieberman, “Shape analysis of brain ventricles using spharm,” in *MMBIA*. 2001, pp. 171–178, IEEE press. [1](#), [2.4](#)
- [12] L. Shen, J. Ford, F. Makedon, and A. Saykin, “Hippocampal shape analysis surface-based representation and classification,” in *SPIE-Medical Imaging*, 2003. [1](#)
- [13] S. Pizer, D. Fritsch, P. Yushkevich, V. Johnson, and E. Chaney, “Segmentation, registration, and measurement of shape variation via image object shape,” *IEEE Trans. Med. Imaging*, vol. 18, pp. 851–865, Oct. 1999. [1](#)
- [14] M. Styner, G. Gerig, J. Lieberman, D. Jones, and D. Weinberger, “Statistical shape analysis of neuroanatomical structures based on medial models,” *Medical Image Analysis*, vol. 7, no. 3, pp. 207–220, 2003. [1](#)
- [15] P. Golland, W.E.L. Grimson, and R. Kikinis, “Statistical shape analysis using fixed topology skeletons: Corpus callosum study,” in *IPMI*, 1999, pp. 382–388. [1](#)
- [16] A. Kelemen, G. Székely, and G. Gerig, “Elastic model-based segmentation of 3d neuroradiological data sets,” *IEEE Trans. Med. Imaging*, vol. 18, pp. 828–839, October 1999. [2.1](#)
- [17] C. Brechbühler, *Description and Analysis of 3-D Shapes by Parametrization of Closed Surfaces*, 1995, Diss., IKT/BIWI, ETH Zürich, ISBN 3-89649-007-9. [2.1](#)
- [18] D. Pantazis, R.M. Leahy, T.E. Nichol, and M. Styner, “Statistical surface-based morphometry using a non-parametric approach,” in *Int. Symposium on Biomedical Imaging(ISBI)*, April 2004, pp. 1283–1286. [2.1](#), [2.6](#)
- [19] M Styner, K Rajamani, LP Nolte, G. Zsemlye, G Szekely, C Taylor, and RH Davies, “Evaluation of 3d correspondence methods for model building,” in *IPMI*, July 2003, pp. 63–75. [2.1](#)
- [20] F.L. Bookstein, *Morphometric Tools for Landmark Data: Geometry and Biology*, Cambridge University Press, 1991. [2.4](#)
- [21] M Styner, JA Lieberman, D Pantazis, and G Gerig, “Boundary and medial shape analysis of the hippocampus in schizophrenia,” *Medical Image Analysis*, vol. 8, no. 3, pp. 197–203, 2004. [2.5](#)
- [22] M Styner, JA Lieberman, RK McClure, DR Weingberger, DW Jones, and G Gerig, “Morphometric analysis of lateral ventricles in schizophrenia and healthy controls regarding genetic and disease-specific factors,” *Proceedings of the National Academy of Science, PNAS*, vol. 102, no. 12, pp. 4872–4877, 2005. [2.5](#)
- [23] Seber GAF, *Multivariate Observations*, John Wiley & Sons, 1984. [2.5](#)
- [24] E. S. Edgington, Ed., *Randomization Tests*, Academic Press, 1995. [2.5](#)
- [25] M. K. Chung, *Statistical Morphometry in Neuroanatomy*, Ph.D. thesis, McGill University, Montreal, 2001. [2.6](#)
- [26] K. J. Worsley, S. Marrett, P. Neelin, A. C. Vandal, K. J. Friston, and A. C. Evans, “A unified statistical approach for determining significant signals in images of cerebral activation,” *Human Brain Mapping*, vol. 4, pp. 58–73, 1996. [2.6](#)

- [27] T.E. Nichols and A.P. Holmes, “Nonparametric analysis of pet functional neuroimaging experiments: A primer with examples,” *Hum Brain Mapp*, vol. 1, no. 15, pp. 1–25, Jan 2001. [2.6](#)
- [28] D. Pantazis, T. E. Nichols, S. Baillet, and R. M. Leahy, “Spatiotemporal localization of significant activation in meg using permutation tests,” in *Information Processing in Medical Imaging*, 2003, pp. 512–523. [2.6](#)
- [29] E. R. Sowell, P. M. Thompson, C. J. Holmes, T. L. Jernigan, and A. W. Toga, “In vivo evidence for post-adolescent frontal and striatal maturation,” *Nature Neuroscience*, Sept./ Oct. 1999. [2.6](#)
- [30] Y Hochberg Y Benjamini, “Controlling false discovery rate: A practical and powerful approach to multiple testing,” *Stat. Soc. Ser. B*, , no. 57, pp. 289–300, 1995. [2.6](#)
- [31] CR Genovese, NA Lazar, and T Nichols, “Thresholding of statistical maps in functional neuroimaging using the false discovery rate,” *NeuroImage*, , no. 15, pp. 870–878, 2002. [2.6](#)
- [32] R. Whitaker, “Applying landmark methods to biological outline data,” in *IEEE Volume Visualization and Graphics Symposium*, 2000, pp. 22–32. [3](#)
- [33] M Styner, M Jomier, and G Gerig, “Closed and open source neuroimage analysis tools and libraries at unc,” in *IEEE Symposium on Biomedical Imaging ISBI*, 2006, pp. 702–705. [3.6](#)
- [34] JJ Levitt, RW McCarley, CC Dickey, MM Voglmaier, MA Niznikiewicz, LJ Seidman, Y Hirayasu, AA Ciszewski, R Kikinis, FA Jolesz, and ME Shenton, “Mri study of caudate nucleus volume and its cognitive correlates in neuroleptic-naive patients with schizotypal personality disorder,” *Am J Psychiatry*, vol. 159, no. 7, pp. 1190–1197, 2002. [4](#)