Hippocampal shape analysis: surface-based representation and classification

Li Shen[a*], James Ford[a], Fillia Makedon[a], and Andrew Saykin[b]

[a]Dartmouth Experimental Visualization Laboratory, Department of Computer Science, Dartmouth College, Hanover, NH 03755, USA
[b]Brain Imaging Laboratory, Departments of Psychiatry and Radiology, Dartmouth Medical School, Lebanon, NH 03756, USA

ABSTRACT

Surface-based representation and classification techniques are studied for hippocampal shape analysis. The goal is twofold: (1) develop a new framework of salient feature extraction and accurate classification for 3D shape data; (2) detect hippocampal abnormalities in schizophrenia using this technique. A fine-scale spherical harmonic expansion is employed to describe a closed 3D surface object. The expansion can then easily be transformed to extract only shape information (*i.e.*, excluding translation, rotation, and scaling) and create a shape descriptor comparable across different individuals. This representation captures shape features and is flexible enough to do shape modeling, identify statistical group differences, and generate similar synthetic shapes. Principal component analysis is used to extract a small number of independent features from high dimensional shape descriptors, and Fisher's linear discriminant is applied for pattern classification. This framework is shown to be able to perform well in distinguishing clear group differences as well as small and noisy group differences using simulated shape data. In addition, the application of this technique to real data indicates that group shape differences exist in hippocampi between healthy controls and schizophrenic patients.

Keywords: Shape analysis, surface modeling, classification, hippocampus, MRI

1. INTRODUCTION

Extracting knowledge about the morphological characteristics of brain structures for medical diagnosis is an important and challenging problem in medical image analysis. One approach is volumetric analysis, which is based on measuring the volume of a structure of interest or of a number of partitions of the structure. The main advantage of volumetric analysis is its simplicity; however, many structural differences may be overlooked. A newer approach, shape analysis, has the potential to provide important information above and beyond simple volume measurements and may characterize abnormalities in the absence of volume differences.

In this paper, surface-based representation and classification techniques are studied for hippocampal shape analysis. The goal is twofold: (1) develop a new approach of salient feature extraction from 3D shapes and accurate classification between 3D shape groups; (2) detect hippocampal abnormalities in schizophrenia using this technique. Object classification involves examples from distinct populations (or *classes*). The aim is to learn a classifier from a training set, or set of labeled examples representing different classes, and then use the classifier to predict the class of any new example with some degree of accuracy. Clearly, an accurate medical classifier has clinical value. This paper focuses on classification for 3D closed surface objects; many brain structures, such as hippocampus, belong to this category and have definable surfaces from which shape information can be

* This work was, in part, supported by NSF IDM 0083423 "Mining Human Brain Data"
Send correspondence to Li Shen. E-mails: {li,jford,makedon}@cs.dartmouth.edu, saykin@dartmouth.edu

extracted. In addition to pick a good classification technique, it is also an important research issue to study how to represent a 3D surface and extract salient features to help make successful classifications.

There are several related previous studies. Csernansky *et al.*[1] studied hippocampal morphometry using a deformation representation by mapping a template image to individuals, and achieved 80% jackknife classification accuracy through a Principal Component Analysis (PCA) and a linear discriminant classification. This approach is image-based, which may not be desirable in more general situations where a surface is not embedded in an image but defined in another way such as segmented boundaries or triangulations. Golland *et al.*[2] did an amygdala-hippocampus complex study, used distance transformation maps as shape descriptors, and achieved 73% accuracy using Support Vector Machine (SVM). We[3] studied hippocampal shape classification, used a symmetrical alignment model with binary image format, and achieved 96% accuracy using only the second principal component after PCA. Note that distance transformation maps and binary images are both voxel-based representations. Their manipulation often involves noisy steps like resampling. Thus, they may not be ideal to represent the shape of a 3D object that can be defined just by its surface.

The SPHARM description[4] is a parametric surface description using spherical harmonics as basis functions, with applications in model-based segmentation[5] and 3D medial shape modeling.[6] It is also used by Gerig and Styner[7] and Shenton *et al.*[8] for shape classification of amygdala-hippocampus complex: they calculated the asymmetry between SPHARM surfaces of left and right complexes, combined it with volume, and achieved 87% accuracy using SVM. The asymmetry measure is a single and highly summarized value for characterizing a configuration of two objects, and so is not a detailed description for 3D objects. However, the SPHARM description itself is a powerful and fine-scale representation for 3D closed surfaces.

With these observations, we propose a new framework for 3D surface object classification that combines SPHARM surface description with a suitable pattern classification technique. A fine-scale spherical harmonic expansion is employed to describe a closed 3D surface object. The expansion can then easily be transformed to extract only shape information (*i.e.*, excluding translation, rotation, and scaling) and create a shape descriptor comparable across different individuals. This representation captures shape features and is flexible enough to do shape modeling, identify statistical group differences, and generate similar synthetic shapes. Principal Component Analysis (PCA) is used to extract a small number of independent features from high dimensional shape descriptors, and Fisher's Linear Discriminant (FLD) is applied for pattern classification. This framework is shown to be able to perform well in distinguishing clear group differences as well as small and noisy group differences using simulated shape data. In addition, the application of this technique to real data indicates that there exists hippocampal shape abnormality in schizophrenia.

The rest of the paper is organized as follows. Section 2 describes the SPHARM expansion for surface shape description and its applications to shape modeling, synthetic shape generation and statistical analysis of group difference. Section 3 presents PCA for feature extraction, FLD for pattern classification, a technique of localizing significant region for better classification, and experimental results. Section 4 concludes the paper.

2. SURFACE-BASED REPRESENTATION

In this study, the test data used to demonstrate our techniques are hippocampus structures extracted from MR (Magnetic Resonance) scans. There are 21 healthy controls and 35 schizophrenic patients involved. The left and right hippocampi in each MR image are identified and segmented by manual tracing in each acquisition slice using the BRAINS software package.[9] A 3D binary image is reconstructed from each set of 2D hippocampus segmentation results, with isotropic voxel values corresponding to whether each voxel is excluded or included. The surface of this 3D binary image is composed of a mesh of square faces (see Figure 1<a-d>(1) for several examples of such surfaces). In this section, we present how to create a shape descriptor for such a closed 3D surface, and show that this shape representation approach is powerful enough for many applications, such as shape modeling, synthetic shape generation, and statistical shape analysis of group differences.

2.1. SPHARM shape description

We adopt SPHARM expansion technique[4] to create a shape description for 3D closed surfaces. An input object surface is assumed to be defined by a square surface mesh converted from an isotropic voxel representation, as

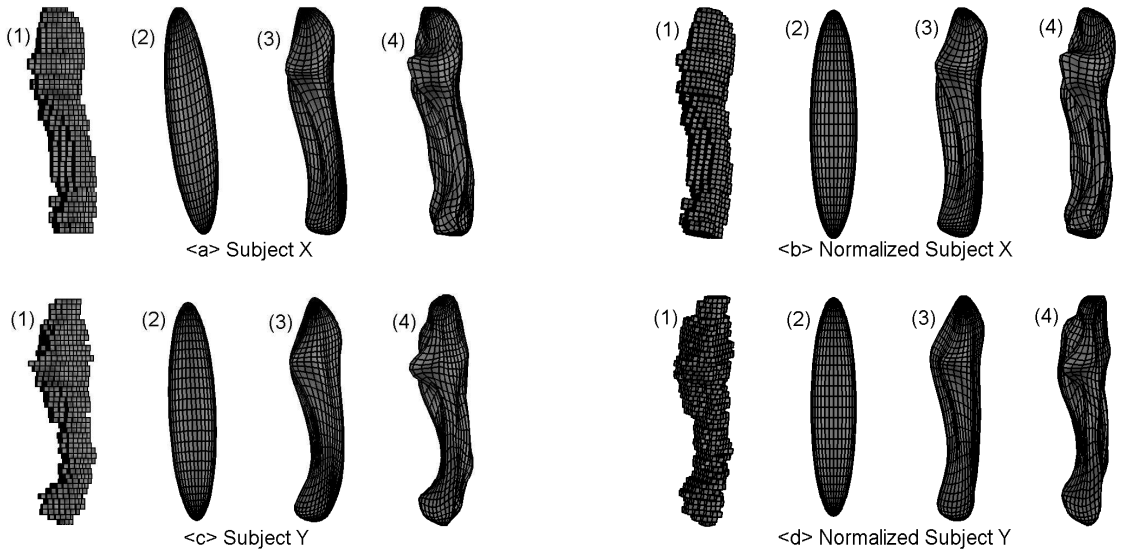

Figure 1. SPHARM description and normalization. <**a,c**> Visualization of left hippocampi of Subject X and Subject Y. <**b,d**> Normalized left hippocampi of Subject X and Subject Y. Inside each of <a-d>: (1) Binary object surface, (2-4) SPHARM reconstructions using coefficients up to degrees 1, 5 and 12, respectively.

mentioned above. Three steps are involved in converting the object surface to its SPHARM shape description: (1) **surface parameterization**, (2) **SPHARM expansion**, and (3) **SPHARM normalization**.

Surface parameterization aims to create a continuous and uniform mapping from the object surface to the surface of a unit sphere. The parameterization is formulated as a constrained optimization problem with the goals of topology and area preservation and distortion minimization (see Brechbhler *et al.*[4] for details). The result is a mapping of two spherical coordinates θ and ϕ to each point $\mathbf{v}(\theta, \phi)$ on a surface:

$$\mathbf{v}(\theta, \phi) = \begin{pmatrix} x(\theta, \phi) \\ y(\theta, \phi) \\ z(\theta, \phi) \end{pmatrix}. \tag{1}$$

When the free variables θ and ϕ run over the whole sphere, $\mathbf{v}(\theta, \phi)$ runs over the whole object surface. **SPHARM expansion** is then used to expand the object surface into a complete set of SPHARM basis functions Y_l^m, where Y_l^m denotes the spherical harmonic of degree l and order m (see Brechbhler *et al.*[4] for details). The expansion takes the form:

$$\mathbf{v}(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \mathbf{c}_l^m Y_l^m(\theta, \phi), \tag{2}$$

where

$$\mathbf{c}_l^m = \begin{pmatrix} c_{xl}^m \\ c_{yl}^m \\ c_{zl}^m \end{pmatrix}. \tag{3}$$

The coefficients $\mathbf{c}_l^m$ are 3D vectors. Their components, c_{xl}^m, c_{yl}^m, and c_{zl}^m are usually complex numbers. The coefficients up to a user-desired degree can be estimated by solving a set of linear equations in a least square fashion. The object surface can be reconstructed using these coefficients, and using more coefficients leads to a more detailed reconstruction. Please see Figure 1<a,c> for two examples. Thus, a set of coefficients actually form an object surface description.

SPHARM normalization is used to create a normalized set of SPHARM coefficients to form a shape descriptor (*i.e.*, excluding translation, rotation, and scaling). The *rotation invariance* is achieved using the

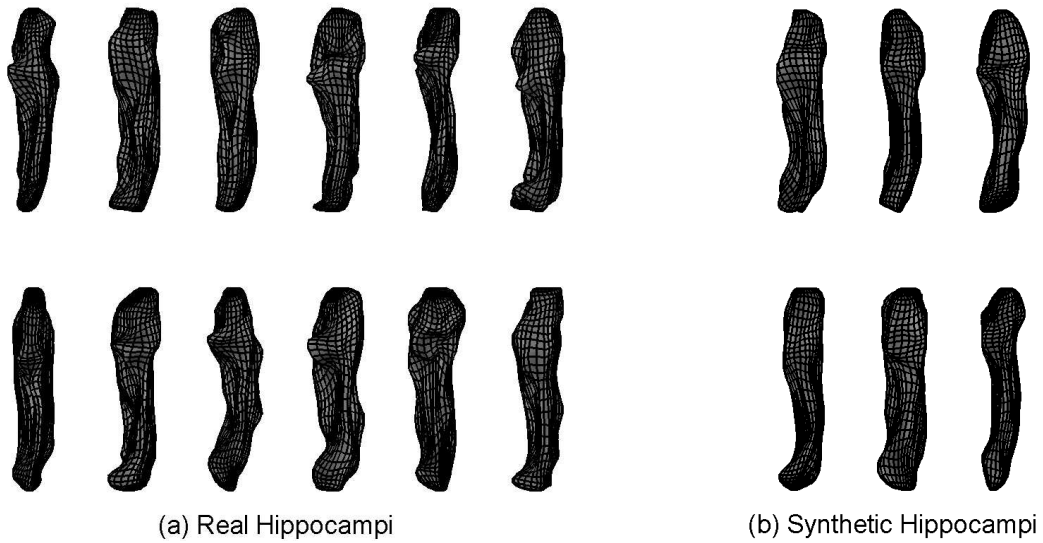

Figure 2. Normalized SPHARM reconstruction of real and synthetic left hippocampi. (a) Left hippocampi of 6 controls (first row) and 6 patients (second row). **(b)** 3 synthetic hippocampi using left control group model (first row) and another 3 using left patient model (second row).

degree 1 reconstruction, which is always an ellipsoid. The parameter net on this ellipsoid is rotated to a canonical position such that the north pole is at one end of the longest main axis, and the crossing point of the zero meridian and the equator is at one end of the shortest main axis. In the object space, the ellipsoid is rotated to make its main axes coincide with the coordinate axes, putting the shortest axis along x and longest along z. The *scaling invariance* can be achieved by dividing all the coefficients by a scaling factor f. In our experiments, we choose f so that the volume of the degree 1 ellipsoid is normalized. Ignoring the degree 0 coefficient results in *translation invariance*. The normalization technique described here works only if a degree 1 ellipsoid is a real ellipsoid, such as with hippocampus data, but not an ellipsoid of revolution or a sphere. In the latter case, higher degree coefficients might need to be involved for the normalization.

After the above steps, a set of canonical coordinates (*i.e.*, normalized coefficients) can be obtained to form a shape descriptor for each surface, and these shape descriptors are comparable across objects. Figure 1 shows two examples of left hippocampi before (see <a,c>) and after (see <b,d>) normalization, where both binary object surfaces and SPHARM reconstructions are presented. Note that, in each of <b,d>, the reconstruction is obtained by using the normalized coefficients (*i.e.*, the shape descriptor), and the binary object surface is transformed accordingly based on the corresponding normalization process. Figure 2(a) shows the normalized SPHARM reconstruction of left hippocampi of 6 controls and 6 patients using their shape descriptors.

For simplicity, *normalized SPHARM coefficients* are hereafter referred to as *SPHARM coefficients*.

2.2. SPHARM coefficients and synthetic data generation

SPHARM coefficients are complex numbers, whose real parts and imaginary parts we treat as separate *elements*. A vector of all these elements forms a shape descriptor for a 3D closed surface. These vectors are related to the same reference system and can be compared across individuals. We assume a normal distribution for each vector element. Given a group of shapes, the mean and standard deviation of each element can be estimated to form a shape model, which may provide some degree of understandings of a shape group. Figure 3 shows sample mean and standard deviation results for 4 different hippocampus groups: left of controls (LC), right of controls (RC), left of patients (LP), and right of patients (RP). A mean vector derives a mean shape. Figure 5(2-5) shows mean shapes of the 4 groups (the surface color will be explained in Section 2.3, please ignore it for now).

The shape model described above can be used to create similar synthetic shapes. For example, a synthetic hippocampus can be constructed if, for each vector element, we draw a random number from its corresponding

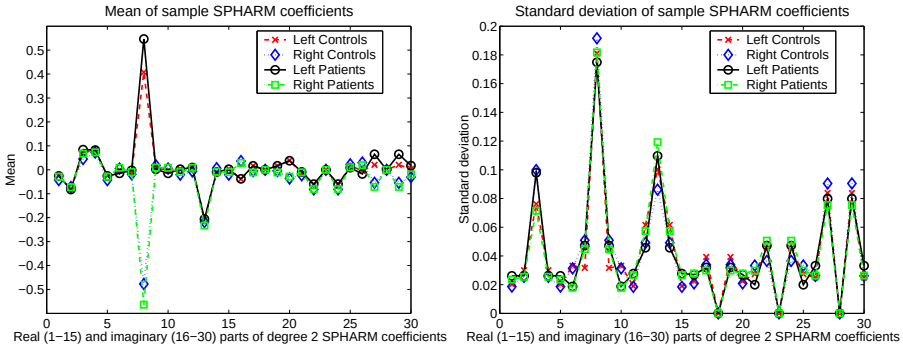

Figure 3. Mean and standard deviation of sample SPHARM coefficients. Real and imaginary parts of degree 2 coefficients are shown for each hippocampus group.

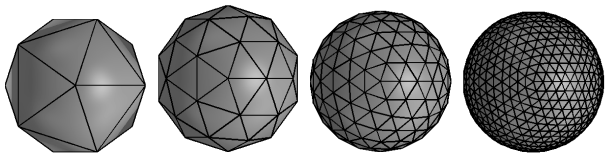

Figure 4. Landmark sampling. Nearly uniform sampling of spherical surfaces by icosahedron subdivision (levels 0-3).

normal distribution with the estimated mean and standard deviation; see Figure 2(b) for 3 synthetic hippocampi using LC shape model and the other 3 using LP model, which look quite similar to real ones in Figure 2(a). Note that this is a very simple approach to synthetic shape generation, where vector elements are assumed independent. However, it is an effective one for our purpose: it can create two groups of shapes that have small and noisy group difference to evaluate our classification approach.

In fact, the SPHARM representation allows for the development of more complicated shape modeling techniques (*e.g.*, Kelemen *et al.*[5]) which can be used for synthetic data generation or even model-based segmentation.

2.3. Landmark dual representation and statistical group difference

It is not easy to intuitively understand a SPHARM coefficient, since the coefficient is usually a complex number and provides a measure of the spatial frequency constituents that compose the object. However, the points of the sampled surface (called *landmarks*) can be considered as a dual representation of the same object. This is a more intuitive descriptor, and so we choose to use this representation in our study. Based on a nearly uniform icosahedron subdivision of spherical surfaces (see Figure 4), we obtain a landmark representation directly from the coefficients via a linear mapping described in Equation 2. Thus, each shape is represented by a set of n landmarks (*i.e.*, sampling points), which are consistent from one shape to the next. In the experiments, we use icosahedron subdivision level 3, and each object has $n = 642$ landmarks.

Now the shape descriptor becomes a $3n$ element vector

$$\mathbf{x} = (x_1, ..., x_n, y_1, ..., y_n, z_1, ..., z_n)^T. \tag{4}$$

Given a group of N shapes, the mean shape $\bar{\mathbf{x}}$ can be calculated using

$$\bar{\mathbf{x}} = \frac{1}{N} \sum_{i=1}^{N} \mathbf{x}_i, \tag{5}$$

where $\mathbf{x}_i$ is the landmark shape descriptor of the i-th shape. The standard deviation can also be estimated for each element in the descriptor. Thus each landmark corresponds to a vector of 3 standard deviation values. The local shape variability at each landmark location can then be measured using the magnitude of the corresponding standard deviation vector: a big magnitude indicates a large local variability.

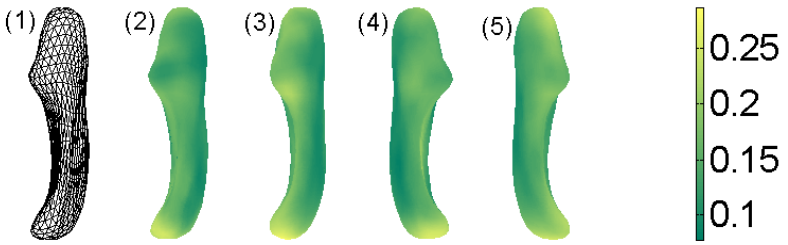

Figure 5. Mean shapes of hippocampus groups. (1) A mean hippocampi visualized by the icosahedron sampling mesh, where mesh vertices correspond to surface landmarks. **(2-5)** Mean shapes of different groups: (2) left controls, (3) left patients, (4) right controls, (5) right patients; where color indicates the magnitude of the standard deviation vector at the corresponding landmark location, and yellow/light (green/dark) color implies more (less) variability at that location.

Figure 5(1) shows the mean shape of the LC group using the icosahedron sampling mesh, where mesh vertices correspond to surface landmarks. Figure 5(2-5) shows the mean shapes of 4 groups: LC, LP, RC and RP, where the surface color indicates the magnitude of the standard deviation vector at the corresponding location, and yellow/light (green/dark) color implies more (less) variability at that location.

The above analysis and visualization provide us a rough shape model for each class. Direct comparison between two models may not derive interesting group differences, since the models contain only summarized information. To find statistical group differences, we do the following.

Given two groups of shapes, a two-sample t-test[10] can be performed for each vector element to obtain an associated p-value. Thus, each landmark location corresponds to a vector of 3 p-values, which represent the probabilities of no group difference along x, y and z directions at that location, respectively. A small p-value indicates a more likely group difference. It may not be good to use the magnitude of the p-value vector to measure an effect combined by all three directions because of the following scenario. Let $V_p = [p_x, p_y, p_z]$ be a p-value vector for a landmark. If the group difference is very obvious along x direction but not y and z directions, we will get a small p_x and large p_y and p_z. As a result, V_p won't have a small magnitude due to the large p_y and p_z and we won't detect actual group difference at this location.

To overcome the above problem, we define the *significance value* (*s-value* in short) as the absolute value of the logarithm of a p-value

$$\text{s-value} = \mid \log(\text{p-value}) \mid . \tag{6}$$

We use s-value to detect the group difference along individual direction, and use the magnitude of the s-value vector to measure an effect combined by all three directions. A large s-value or vector magnitude implies that group differences are more likely to exist. Figure 6 shows the mean shape of all left hippocampi and the mean shape of all right ones in two different views, where the color indicates either the magnitude of the corresponding s-value vector (in subplots labeled (1)) or the s-values along x, y and z directions (in subplots labeled (2-4) respectively). Yellow/light (green/dark) color implies that group difference is more (less) likely to exist there. By using s-values in Figure 6, we can see that the subplot (1) becomes a good combination of the corresponding subplots (2-4) in terms of the yellow/light color area, which is the area showing significant group differences. Thus it seems enough to just use the magnitude of the s-value vector for visualizing statistical group differences. This visualization (Figure 6) shows that the group difference appears in the head/anterior region as well as the tail/posterior region for both left and right hippocampi.

3. PATTERN CLASSIFICATION

We have shown that the SPHARM description and its dual landmark representation are powerful enough for many applications such as shape modeling, synthetic shape simulation, and statistical analysis of group differences. However, our ultimate goal is to build a model for accurate classification of new examples. This is a more difficult task, especially when we use a fine-scale representation involving many landmarks. In this section, we

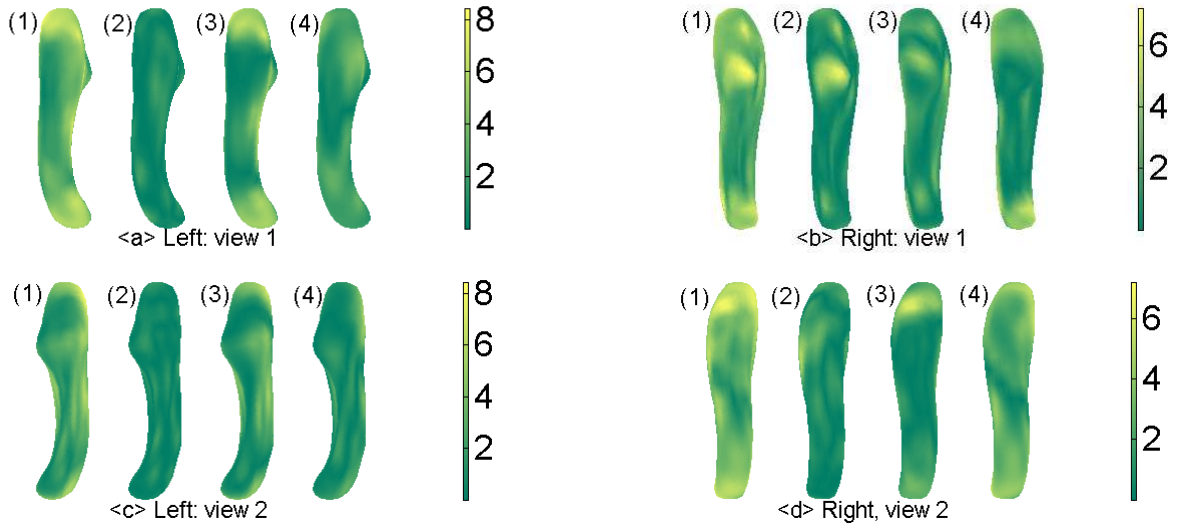

Figure 6. Statistical group differences. <**a,c**> Two views of the mean of all left hippocampi. <**b,d**> Two views of the mean of all right hippocampi. Inside each of <a-d>: Color indicates (1) the magnitude of the corresponding s-value vector, (2-4) the s-values along x, y and z directions respectively. Yellow/light (green/dark) color implies that group difference is more (less) likely to be present at that location.

present techniques for dimensionality reduction and pattern classification, provide experimental results to show the effectiveness of our approach, and discuss how to localize significant area for a better classification.

3.1. Dimensionality reduction

We use a fine-scale landmark representation for classification, where each shape is described by 642 landmarks and the shape descriptor has a length of $642 * 3 = 1926$. Principal Component Analysis (PCA)[11] is applied to reduce dimensionality. This involves eigenanalysis of the covariance matrix $\mathbf{\Sigma}$ of the data as follows

$$\mathbf{\Sigma} = \frac{1}{N-1}\sum_{i=1}^{N}(\mathbf{x_i} - \bar{\mathbf{x}})(\mathbf{x_i} - \bar{\mathbf{x}})^T, \tag{7}$$

$$\mathbf{\Sigma P} = \mathbf{DP}, \tag{8}$$

where the columns of $\mathbf{P}$ hold eigenvectors, and the diagonal matrix $\mathbf{D}$ holds eigenvalues of $\mathbf{\Sigma}$. The eigenvectors in $\mathbf{P}$ can be ordered according to respective eigenvalues, which tell the amount of variance explained by each eigenvector. The first few eigenvectors (with greatest eigenvalues) often explain most of variance in the data. Now any shape $\mathbf{x}$ in the data can be obtained using

$$\mathbf{x} = \bar{\mathbf{x}} + \mathbf{Pb}, \tag{9}$$

where $\mathbf{b}$ is a vector containing the components of $\mathbf{x}$ in basis $\mathbf{P}$, which are called *principal components*. Since eigenvectors are orthogonal, $\mathbf{b}$ can be obtained using

$$\mathbf{b} = \mathbf{P}^T(\mathbf{x} - \bar{\mathbf{x}}). \tag{10}$$

Assume a dataset contains m objects and so its inherent dimensionality is $\leq m-1$. It can be shown that the first $m-1$ principal components are enough to capture all the variance in the dataset. The PCA basis $\mathbf{P}$ can be formed by just the first $m-1$ eigenvectors and then Equations 9 and 10 still hold. In this case, $\mathbf{b}$ becomes an $m-1$ element vector, which can be thought of a new and more compact representation of the shape $\mathbf{x}$ in the new basis of the deformation modes. Due to this compactness, we use $\mathbf{b}$ as the new descriptor in the classification. In addition, the first few components in $\mathbf{b}$ can approximate $\mathbf{x}$, if the respective eigenvectors explain a large amount of variance in the data. Hereafter, $\mathbf{b}$ is referred to as a *feature vector*, and eigenvectors as *PC dimensions*.

3.2. Classifier design

We use Fisher's Linear Discriminant (FLD) as our classifier. FLD is a multi-class technique for pattern classification. FLD projects a training set (consisting of c classes) onto $c-1$ dimensions such that the ratio of between-class and within-class variability is maximized, which occurs when the FLD projection places different classes into distinct and tight clumps.[11]

This optimal projection $\mathbf{W}_{\text{opt}}$ is calculated as follows. Assume that we have a set of n d-dimensional samples $\mathbf{x}_1, ..., \mathbf{x}_n$, n_i in the subset $\mathcal{D}_i$ labeled ω_i, where $n = \sum_{k=1}^{c} n_k$ and $i \in \{1, ..., c\}$. We define the between-class scatter matrix $\mathbf{S}_B$ and the within-class scatter matrix $\mathbf{S}_W$ as

$$\mathbf{S}_B = \sum_{i=1}^{c} |\mathcal{D}_i|(\mathbf{m}_i - \mathbf{m})(\mathbf{m}_i - \mathbf{m})^T, \tag{11}$$

$$\mathbf{S}_W = \sum_{i=1}^{c} \sum_{\mathbf{x} \in \mathcal{D}_i} (\mathbf{x} - \mathbf{m}_i)(\mathbf{x} - \mathbf{m}_i)^T, \tag{12}$$

where $\mathbf{m}$ is the mean of all samples and $\mathbf{m}_i$ the mean of class ω_i. If $\mathbf{S}_W$ is nonsingular, the optimal projection $\mathbf{W}_{\text{opt}}$ is chosen by

$$\mathbf{W}_{\text{opt}} = \underset{\mathbf{W}}{\text{argmax}} \frac{|\mathbf{W}^T \mathbf{S}_B \mathbf{W}|}{|\mathbf{W}^T \mathbf{S}_W \mathbf{W}|} = [\mathbf{w}_1 \mathbf{w}_2 \dots \mathbf{w}_m] \tag{13}$$

where $\{\mathbf{w}_i \mid i = 1, 2, \dots, m\}$ is the set of generalized eigenvectors of $\mathbf{S}_B$ and $\mathbf{S}_W$ corresponding to set of decreasing generalized eigenvalues $\{\lambda_i \mid i = 1, 2, \dots, m\}$, *i.e.*,

$$\mathbf{S}_B \mathbf{w}_i = \lambda_i \mathbf{S}_W \mathbf{w}_i. \tag{14}$$

Note that an upper bound on m is $c-1$. See Duda *et al.*[11] for details.

In our case, we have only two classes, and so the above FLD basis $\mathbf{W}_{\text{opt}}$ becomes just a column vector $\mathbf{w}$. Once this $\mathbf{w}$ has been found, a new feature vector can be projected onto $\mathbf{w}$ to classify it. The resulting scalar value can be compared to the projections of the training set on the same basis $\mathbf{w}$. In this one dimensional FLD space, we propose the following Bayesian approach to perform classification.

A normal distribution $\mathcal{N}(\mu_i, \sigma_i^2)$ is assumed in the FLD space for each class ω_i, where its mean μ_i and variance σ_i^2 can be estimated from the training set. The conditional probability $P(\mathbf{x} \in \mathcal{D}_i \mid y)$ that a new object $\mathbf{x}$ belongs to class ω_i (*i.e.*, the label of $\mathcal{D}_i$), conditioned on its FLD projection being y, can then be calculated by the following Bayesian model[11]

$$\begin{aligned} P(\mathbf{x} \in \mathcal{D}_i \mid y) &= \frac{p(y \mid \mathbf{x} \in \mathcal{D}_i) * P(\mathbf{x} \in \mathcal{D}_i)}{\sum_{j=1}^{c} p(y \mid \mathbf{x} \in \mathcal{D}_j) * P(\mathbf{x} \in \mathcal{D}_j)} \\ &= \frac{p(y \mid y \sim \mathcal{N}(\mu_i, \sigma_i^2)) * P(\mathbf{x} \in \mathcal{D}_i)}{\sum_{j=1}^{c} p(y \mid y \sim \mathcal{N}(\mu_j, \sigma_j^2)) * P(\mathbf{x} \in \mathcal{D}_j)}. \end{aligned} \tag{15}$$

In Equation 15, $p(y \mid \mathbf{x} \in \mathcal{D}_i)$ is the state-conditional Probability Density Function (pdf) for random variable y conditioned on $\mathbf{x}$ belonging to class ω_i, which is equivalent to pdf $p(y \mid y \sim \mathcal{N}(\mu_i, \sigma_i^2))$ based on our assumption of a normal distribution for y. The prior probability $P(\mathbf{x} \in \mathcal{D}_i)$ of $\mathbf{x}$ belonging to class ω_i are chosen as the fraction of the dataset belonging to ω_i. Our approach assigns a new object to the class corresponding to the largest posterior probability computed by Equation 15. In the equal case, the new object joins the class having the closest mean.

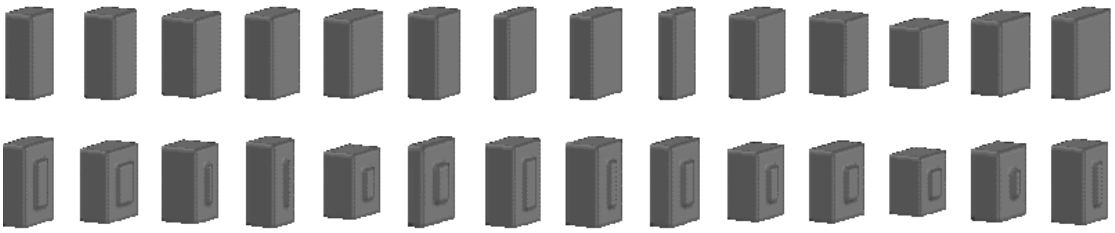

Figure 7. Two classes of synthetic surfaces: The first row shows 14 surfaces in Class 1 (or Cuboid Class), while the second shows 14 surfaces in Class 2 (or Cuboid-bump Class).

3.3. Experimental results

Classification is tested using a **jackknife** methodology (Duda *et al.*[11]): each object is removed in turn as the test case, the remaining objects forms a training set for classifier learning, the resulting classifier is tested on the removed object, and the accuracy is estimated as the mean of these leave-one-out accuracies. Only the training objects are involved in creating the PCA basis $\mathbf{P}$ (derived by Equation 8) and the FLD basis $\mathbf{W}_{\text{opt}}$ (derived by Equation 13). The test object is first projected to $\mathbf{P}$ according to Equation 10, and then projected to $\mathbf{W}_{\text{opt}}$ to get a single FLD value that can be classified using our Bayesian model.

Classification can be performed using just the first few Principal Components (PCs), which account for significant amount of data variance, based on the assumption that this information is crucial for classification and the rest noisy. To determine how many PCs to use, we simply apply a classifier to each data set m times and pick the best result: at each time $i \in \{1, 2, \ldots, m\}$, we use the first i PCs of a feature vector for classification. Given n training objects, we can only have $m = n - 2$ for the FLD classifier. This is because the FLD projection computed by Equation 13 requires a nonsingular $\mathbf{S}_W$ and Equation 12 implies that the rank of $\mathbf{S}_W$ is $\leq n - c$, where n is the training set size and c the number of classes.

Figure 8 shows classification results for 3 experiments, where the jackknife classification accuracy is displayed in (a) and the average variance explained by the first n PCs is plotted in (b). Each trial has its own PCA basis due to the jackknife methodology, and so the variance shown in Figure 8(b) is the average of all trials.

In the first experiment, we have two classes of surfaces (see Figure 7): (1) 14 cuboid surfaces; and (2) 14 cuboid-bump surfaces, each with a centered bump on one face. Although the group shape difference is clear in this example, the intra-group variabilities are also present in each group. The 100% jackknife accuracy is very consistent if we use more than two PCs (corresponding to $> 98\%$ of the variance) and this verifies our technique can effectively detect the group difference in the presence of noisy intra-group variabilities.

The second and third experiments are for synthetic left and right hippocampi respectively: between one group simulated using the control model and the other using the patient model. There are 28 synthetic hippocampi in each group. Please see Section 2.2 for the synthetic shape generation approach. In the left case, $89\% \sim 96\%$ accuracy is achieved when $12 \sim 48$ PCs (corresponding to $80\% \sim 99\%$ of the variance) are used; in the right case, $93\% \sim 100\%$ accuracy is achieved when $18 \sim 53$ PCs (corresponding to $86\% \sim 100\%$ of the variance) are used. Synthetic hippocampi generated from different groups have systematic small differences in terms of mean shape and coefficient variance (see Figures 3 and 2(b)). Our technique seems to be able to capture these differences.

The above three experiments show that our classification framework performs pretty well in distinguishing clear group differences as well as small and noisy group differences.

Figure 9 shows the classification result for the real data set, which includes 21 healthy controls and 35 schizophrenic patients. Two experiments are performed: one for left hippocampi, the other for right hippocampi, both for classification between the control group and the patient group. In most cases, better than chance accuracy can be achieved. For left hippocampi, the best accuracy is 66% when we use $2 \sim 4$ PCs (corresponding to $49\% \sim 63\%$ of the variance); for right hippocampi, $68\% \sim 71\%$ accuracy is achieved when $1 \sim 2$ or $29 \sim 31$ PCs (corresponding to $30\% \sim 42\%$ or $94\% \sim 95\%$ of the variance) are used. These results suggest that there may

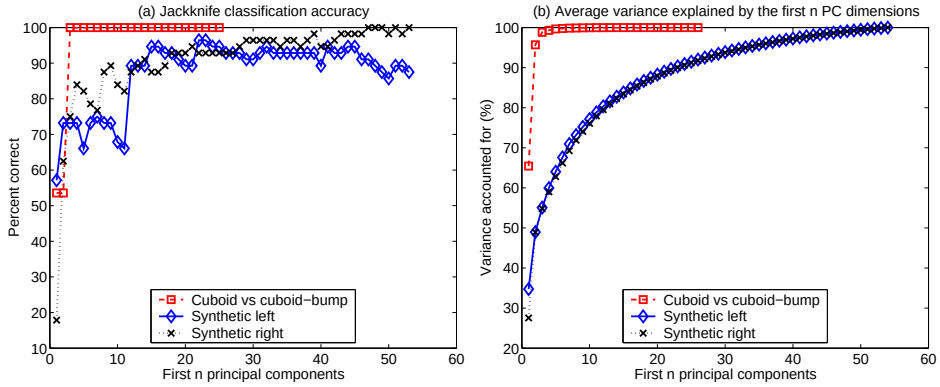

Figure 8. Classification for simulated data. (a) Jackknife classification accuracy. **(b)** Average variance explained by the first n PC dimensions. Three experiments are shown: (1) 14 cuboid surfaces vs. 14 cuboid-bump surfaces, see Figure 7, (2) 28 synthetic left hippocampi simulated using the left control shape model vs. 28 using the left patient model, (3) 28 synthetic right hippocampi simulated using the right control model vs. 28 using the right patient model.

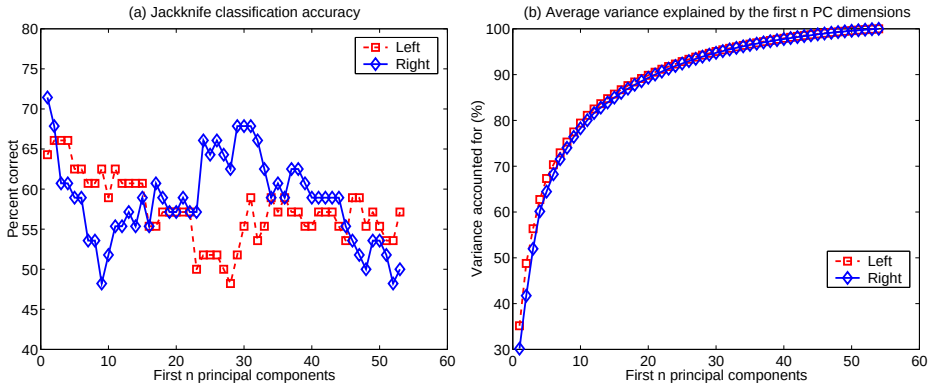

Figure 9. Classification for real hippocampus data. (a) Jackknife classification accuracy. **(b)** Average variance explained by the first n PC dimensions. There are 21 controls and 35 patients involved. Two experiments are shown: (1) controls versus patients for left hippocampi, (2) controls versus patients for right hippocampi.

exist hippocampal shape difference in schizophrenia. It becomes an interesting problem to localize the difference and hopefully further improve the classification accuracy. Section 3.4 addresses this issue.

3.4. Significant region

In the landmark representation, we feel that some landmarks might be more *significant* than others for classification. We call a set of these landmarks a *significant region*. Discovery of such a region may help obtain a better classification. Work is still in progress to study this problem. Here we just present some encouraging preliminary results regarding this research issue.

Please recall that we have done some statistical analysis of group shape differences in Section 2.3 and the result is a vector of s-values for each landmark. The larger the vector magnitude is, the more likely the group difference is to appear there, see Figure 6, and hopefully by intuition the more the corresponding landmark contributes to classification. With this observation, a simple approach to the significant region selection is just to pick the top n landmarks according to the magnitude of the corresponding s-value vector. Figure 10 shows sample significant regions determined using this approach, where 60 out of 642 landmarks are picked to form the region in each of left and right cases.

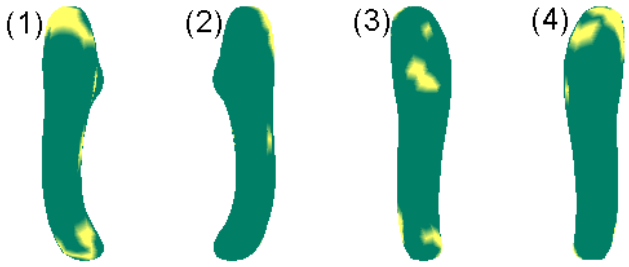

Figure 10. Significant region selection. **(1,2)** Two views of the mean left hippocampus. **(3,4)** Two views of the mean right hippocampus. Yellow/light color indicates the significant region, which is determined by the top 60 (out of 642) landmarks according to the magnitude of the corresponding s-value vector. Please refer to Figure 6 for the visualization of s-values.

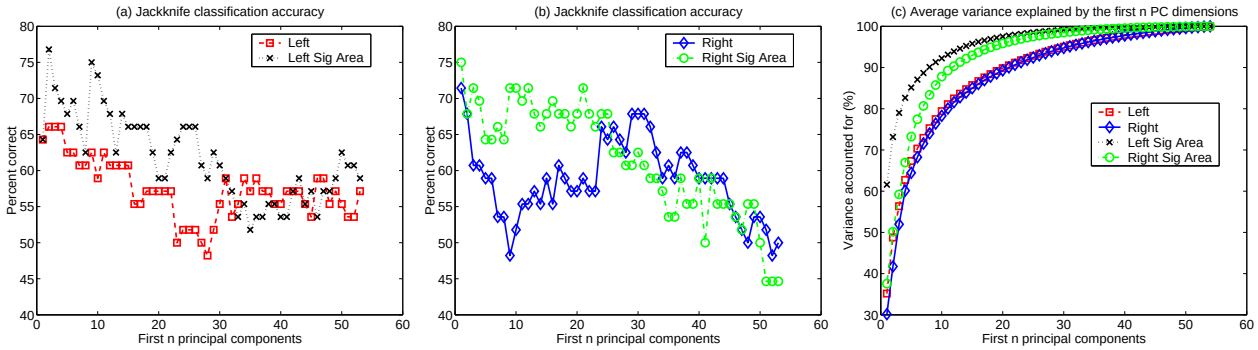

Figure 11. Classification using the significant region versus all the region. (a) Jackknife classification accuracy for left hippocampi. **(b)** Jackknife classification accuracy for right hippocampi. **(c)** Average variance explained by the first n PC dimensions. Four experiments are shown: (1) controls versus patients for left hippocampi using all the region, (2) controls versus patients for left hippocampi using the significant region, (3) controls versus patients for right hippocampi using all the region, (4) controls versus patients for right hippocampi using the significant region.

Classification can be done using just these significant landmarks instead of all. Figure 11 shows a performance comparison between them, where the Significant Regions (SRs) defined in Figure 10 are used in the experiments. A better accuracy can generally be obtained if we use just the first few PCs: in the left case, $1 \sim 32$ PCs (corresponding to $62\% \sim 99\%$ of the variance for the SR data), with an improved best accuracy of 77%; in the right case, $1 \sim 25$ PCs (corresponding to $38\% \sim 98\%$ of the variance for the SR data), with an improved best accuracy of 75%. This result seems reasonable: Figure 11(c) shows that more data variance is captured by the first few PCs in the SF cases than that in the regular cases, and this information becomes enough and important for classification.

4. CONCLUSIONS

Previous studies on hippocampus-related shape classification have employed volumetric representations such as image-based deformation,[1] binary image,[3] and distance map.[2] We observe that, for a 3D volumetric object, its boundary or surface actually defines the shape. Surface-based representation may be more appropriate for shape study unless the appearance or tissue inside the object is also the focus of interest. We have studied surfaced-based representation and classification techniques for hippocampal shape analysis in this paper.

Our shape representation is based on the SPHARM description[4] which has been applied successfully in model-based segmentation[5] and 3D medial shape modeling.[6] It is also used in a shape classification study[7,8] to obtain hippocampal asymmetry measure, which is a single and highly summarized value. We employ the SPHARM description in a much more detailed level, and focus on left and right hippocampi themselves. A closed 3D surface object is described by a fine-scale spherical harmonic expansion, which can then be transformed to extract only

shape information and create a shape descriptor comparable across different individuals. We show that this shape representation is flexible enough to do shape modeling, identify statistical group differences, and generate similar synthetic shapes.

The shape representation is then combined with a suitable classification approach, which includes principal component analysis for feature extraction and Fisher's linear discriminant for pattern classification. This framework is shown to be able to perform well (achieve 96% $\sim$ 100% jackknife accuracy) in distinguishing clear group differences as well as small and noisy group differences using simulated shape data. In addition, the application of this technique to real data achieves 66% accuracy for the left hippocampi and 71% for the right ones and suggests that there may exist hippocampal shape abnormality in schizophrenia. We also present some encouraging preliminary result on how to identify the significant region on the surface for a better classification, and the accuracy is improved to 77% for the left hippocampi and 75% for the right ones. Work is in progress to study other approaches that locate significant regions and hopefully further improve the classification accuracy for real hippocampi.

ACKNOWLEDGMENTS

We thank NSF IDM 0083423, NARSAD, NH Hospital and Ira DeCamp Foundation for support. We are grateful to Hany Farid for helpful discussions, and to Annette Donnelly, Laura A. Flashman, Tara L. McHugh and Molly B. Sparling for creating hippocampal traces.

REFERENCES

1. J. G. Csernansky, S. Joshi, L. Wang, J. W. Halleri, M. Gado, J. P. Miller, U. Grenander, and M. I. Miller, "Hippocampal morphometry in schizophrenia by high dimensional brain mapping," *Proc. Natl. Acad. Sci. USA* **95**, pp. 11406–11411, September, 1998.
2. P. Golland, W. E. L. Grimson, M. E. Shenton, and R. Kikinis, "Small sample size learning for shape analysis of anatomical structures," in *3rd International Conference on Medical Image Computing and Computer Assisted Intervention*, *LNCS* **1935**, pp. 72–82, (Pittsburgh, Pennsylvania, USA), October 11-14, 2000.
3. A. J. Saykin, L. A. Flashman, T. McHugh, C. Pietras, T. W. McAllister, A. C. Mamourian, R. Vidaver, L. Shen, J. C. Ford, L. Wang, and F. Makedon, "Principal components analysis of hippocampal shape in schizophrenia," in *International Congress on Schizophrenia Research*, (Colorado Springs, Colorado, USA), March 29 - April 2, 2003.
4. C. Brechbhler, G. Gerig, and O. Kubler, "Parametrization of closed surfaces for 3d shape description," *Computer Vision and Image Understanding* **61(2)**, pp. 154–170, 1995.
5. A. Kelemen, G. Szekely, and G. Gerig, "Elastic model-based segmentation of 3-d neuroradiological data sets," *IEEE Transactions on Medical Imaging* **18**, pp. 828–839, 1999.
6. M. Styner and G. Gerig, "Three-dimensional medial shape representation incorporating object variability," in *Proc. of Computer Vision and Pattern Recognition II*, pp. 651–656, (Kauai, Hawaii, USA), October 11-14, 2001.
7. G. Gerig and M. Styner, "Shape versus size: Improved understanding of the morphology of brain structures," in *4th International Conference on Medical Image Computing and Computer Assisted Intervention*, *LNCS* **2208**, pp. 24–32, (Utrecht, The Netherlands), October 14-17, 2001.
8. M. Shenton, G. Gerig, R. W. McCarley, G. Szekely, and R. Kikinis, "Amygdala-hippocampal shape differences in schizophrenia: the application of 3d shape models to volumetric mr data," *Psychiatry Research-Neuroimaging* **115**, pp. 15–35, August 20, 2002.
9. Iowa MHCRC Image Processing Lab, *Brains Software*, http://moniz.psychiatry.uiowa.edu.
10. NIST/SEMATECHR, *e-Handbook of Statistical Methods*, http://www.itl.nist.gov/div898/handbook/, 2002.
11. R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification (2nd ed)*, Wiley, New York, NY, 2000.