

Effective Classification of 3D Closed Surfaces: Application to Modeling Neuroanatomical Structures*

Li Shen¹, James Ford¹, Fillia Makedon¹, Andrew Saykin²

¹Dartmouth Experimental Visualization Lab, Computer Science, Dartmouth College, Hanover, NH 03755

²Brain Imaging Lab, Psychiatry and Radiology, Dartmouth Medical School, Lebanon, NH 03756

Abstract

We present a new technique for 3D surface object classification that combines a powerful shape representation approach with suitable pattern classification techniques. Spherical harmonic parameterization and normalization techniques are used to describe a surface shape and derive a dual high dimensional landmark representation. A point distribution model is applied to reduce the dimensionality. Linear discriminants and support vector machines are used for classification, with several feature selection approaches. We demonstrate the effectiveness of this technique using a hippocampus dataset and achieve a jackknife classification accuracy of 93%.

1 Introduction

Object classification via shape analysis is an important and challenging problem in computer vision and medical image analysis. In the brain imaging domain, the goal is to identify shape abnormalities in a structure of interest that are associated with a particular condition to aid diagnosis and treatment. This paper considers only simply connected 3D objects, a category many brain structures belong to. We demonstrate our techniques on the hippocampi of 35 schizophrenics and 21 controls.

There are several hippocampus-related shape classification studies. Csernansky et al. [3] studied hippocampal morphometry, using an image-based deformation representation, and achieved 80% jackknife classification accuracy through principal component analysis (PCA) and a linear discriminant. Timoner, Golland, et al. [12] did amygdala-hippocampus complex studies, using distance transformation maps and displacement fields as shape descriptors, and achieved best accuracies of 77% and 87%, respectively, using support vector machines (SVMs). We studied hippocampal shape classification in [11], using a symmetric alignment model and binary images, and achieved 96% accuracy using only the second principal component after PCA.

The above are all image-based or voxel-based approaches. We are more interested in surface-based approaches, which have the following advantages: (1) they capture enough information for shape studies; (2) they

can be applied in more general situations where a surface is not embedded in an image but defined in another way such as a triangulation; and (3) some noisy steps like resampling in voxel-based analyses can be avoided.

The SPHARM description [1] is a parametric surface description using spherical harmonics as basis functions, with applications in model-based segmentation [7] and 3D medial shape (m-rep) modeling [10]. Gerig, Styner, et al., have done numerous studies (see [6] for a list) identifying statistical shape abnormalities of different neuroanatomical structures using SPHARM and m-rep. They have used SPHARM in [5] for calculating hippocampal asymmetry, combined it with volume, and achieved 87% accuracy using SVM. We have also used the SPHARM model together with PCA and Fisher's linear discriminant (FLD) in a hippocampal shape classification study [9] and achieved 77% accuracy. This paper extends our previous work by combining SPHARM, the point distribution model (PDM) [2], and FLD and SVM with effective feature selection to obtain an improved accuracy of 93%. In addition, we introduce a new method for visualizing discriminative patterns.

2 Shape Description

We adopt the SPHARM expansion technique [1] 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 (see Figure 1(a) for a hippocampal surface). Three steps are involved to obtain a 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. The result is a bijective mapping between each point \mathbf{v} on a surface and spherical coordinates θ and ϕ : $\mathbf{v}(\theta, \phi) = (x(\theta, \phi), y(\theta, \phi), z(\theta, \phi))^T$.

SPHARM expansion expands 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 . The expansion takes the form:

*This work is supported by NSF IDM 0083423, NARSAD, NH Hospital and Ira DeCamp Foundation.




Figure 1: (a) Object surface and SPHARM reconstructions using coefficients up to degrees 1, 5 and 12. (b) Two sample reconstructions: a control and a patient.

$\mathbf{v}(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l \mathbf{c}_l^m Y_l^m(\theta, \phi)$, where $\mathbf{c}_l^m = (c_{xl}^m, c_{yl}^m, c_{zl}^m)^T$. The coefficients \mathbf{c}_l^m up to a user-desired degree can be estimated by solving a set of linear equations in a least squares fashion. The object surface can be reconstructed using these coefficients, and using more coefficients leads to a more detailed reconstruction. See Figure 1(a) for an example.

SPHARM normalization creates a shape descriptor (i.e., excluding translation, rotation, and scaling) from a normalized set of SPHARM coefficients, which are comparable across objects. *Rotation invariance* is achieved by aligning the degree 1 reconstruction, which is always an ellipsoid. *Scaling invariance* can be achieved by dividing all the coefficients by a scaling factor f . In our experiments, we choose f so that the object volume is normalized. Ignoring the degree 0 coefficient results in *translation invariance*. Figure 1(b) shows two reconstruction examples based on this approach.

Using a uniform icosahedron subdivision of spherical surfaces (see [7]), we obtain a dual landmark representation from the coefficients via the linear mapping described in SPHARM expansion formula. This is a more intuitive descriptor than coefficients, and so we choose to use this representation. In our experiments we use icosahedron subdivision level 3, resulting in $n = 642$ landmarks for each object and $3n = 1926$ feature elements.

Clearly, we have many more dimensions than training objects. PCA [4] is applied to reduce dimensionality to make classification feasible. This involves eigenanalysis of the covariance matrix Σ of the data as follows: $\Sigma = \frac{1}{N-1} \sum_{i=1}^N (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T$, $\Sigma \mathbf{P} = \mathbf{D} \mathbf{P}$, where the columns of \mathbf{P} hold eigenvectors, and the diagonal matrix \mathbf{D} holds eigenvalues of Σ . The eigenvectors in \mathbf{P} can be ordered according to respective eigenvalues, which are proportional to the 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 $\mathbf{x} = \bar{\mathbf{x}} + \mathbf{P}\mathbf{b}$, 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 $\mathbf{b} = \mathbf{P}^T(\mathbf{x} - \bar{\mathbf{x}})$.

Given a dataset of m objects, the first $m - 1$ principal components are enough to capture all the data variance. In this case, \mathbf{b} becomes an $m - 1$ element vector, which can be thought of as a new and more compact representation of the shape \mathbf{x} in the new basis of the deformation modes. Hereafter, \mathbf{b} is referred to as a *feature vector*.

The model described above is a point distribution model (PDM) [2]. PDM is applied to our hippocampal data and generates 2 sets of feature vectors: the **Left Set** for left hippocampi, and the **Right Set** for right ones.

3 Classification

In this section we describe classifiers, discuss feature selection, and give experimental results.

3.1 Fisher’s Linear Discriminant

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. Please refer to [4] for how to calculate this optimal projection \mathbf{W}_{opt} .

Since we have only two classes, the FLD basis \mathbf{W}_{opt} becomes a column vector \mathbf{w} . New feature vectors can be compared to the training set, and thus classified, by projecting them onto \mathbf{w} . We choose four approaches to perform classification in the FLD space: (1) FLD-BM, (2) FLD-1NN, (3) FLD-3NN, and (4) FLD-NM.

FLD-BM assumes a normal distribution in the FLD space for each class. Using a Bayesian model (BM), the certainty that a test subject could be explained by each class’s distribution can be calculated based on the training set (see [9] for details). FLD-BM assigns a new object to the class corresponding to the largest posterior probability computed by the BM.

FLD-1NN and **FLD-3NN** are two k nearest neighbour (kNN) classifiers with $k = 1$ and $k = 3$ respectively. A kNN classifier assigns a new object to the most common class in the k nearest labelled training objects. **FLD-NM** is a nearest mean (NM) classifier, which assigns a new object to the class having the nearest mean.

3.2 Support Vector Machine

A linear classifier corresponds to a decision hyperplane to separate different classes. The margin is defined as the distance from the decision hyperplane to the closest training set exemplar. The aim in training a linear SVM is to find the separating hyperplane with the largest margin; the expectation is that the larger the margin, the better the generalization of the classifier.

We employ the publicly available *OSU SVM Classifier Matlab Toolbox version 3.00* [8]. We denote a linear SVM classifier with the cost of the constraint violation x as **SVM-Cx**. **SVM-C1**, **SVM-C10**, **SVM-C100** are applied in our experiments. We refer the readers to [4, 8]




Figure 2: (a) Two of the synthetic surfaces in class 1 (top) and class 2 (bottom). (b) Third PC of surfaces, which discriminates classes. (c) The mean surface with colors indicating magnitude of contribution to the third PC.

for more technical and implementation details.

3.3 Experimental Results

This work differs from our previous work [9] in three respects: (1) We apply PCA to all data in a single step, rather than constructing a new basis for each jackknife trial based on individual training sets. This is a simpler approach that minimizes representation errors. (2) We add FLD-1NN, FLD-3NN, FLD-NM, SVM-C1, SVM-C10, and SVM-C100 to the FLD-BM classifier. (3) Instead of only using varying numbers of principal components to test classification, in this work we also examine using varying numbers of significant features (SFs).

SFs are based on the observation that selecting an appropriate subset of features often improves classification accuracy by reducing the number of parameters that need to be estimated. In our study, features are principal components (PCs), and we feel that some PCs are more *important* than others for classification, but not necessarily matching the ordering of the variance amounts they explain. Figure 2 shows an example based on two classes of 14 synthetic rectangular surfaces each, with bumps centered on one face in the second class. We apply our classifiers to this data, and the jackknife accuracies using the first i PCs are: $< 60\%$ for $i = 1, 2$; and 100% for $i = 3, \dots, 24$. As shown in Figure 2(b), the 3rd PC alone supports perfect classification, and thus should be considered more important than components that do not. Figure 2(c) shows the contributions of landmark locations to the 3rd PC. From this visualization, it is apparent that the 3rd component is focused on the most significant surface region for discriminating the synthetic classes.

To rank the effectiveness of features, we employ a simple two-sample t-test on each feature and obtain a p-value associated with the test statistic. A lower p-value implies more significant group difference statistically. We say that feature A is more *significant* than feature B if A has a smaller p-value. We hypothesize that more significant features can help more in distinguishing two groups for classification.

We perform 3 sets of experiments in terms of feature selection: (1) **PC**: use the first n PCs; (2) **SF**: use the first

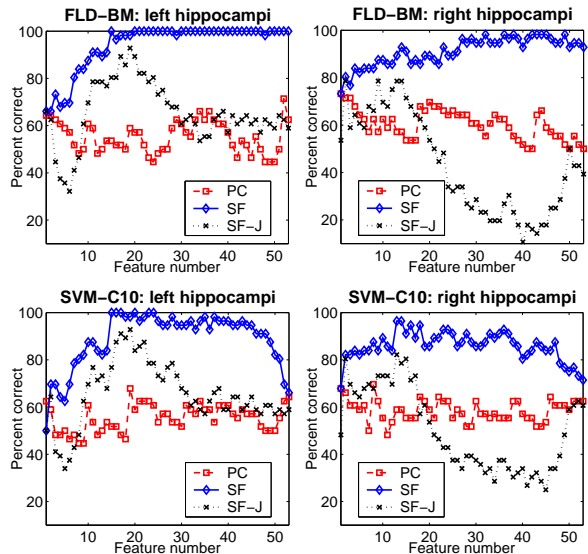


Figure 3: Jackknife classification accuracy.

acc, fts	Left	Right
FLD-BM	93%, 19	79%, 2
FLD-1NN	89%, 17	79%, 11
FLD-3NN	89%, 19	79%, 2
FLD-NM	91%, 19	79%, 14
SVM-C1	80%, 19	82%, 13
SVM-C10	93%, 19	82%, 13
SVM-C100	91%, 17	82%, 15

Table 1: Best jackknife accuracy using **SF-J**: Each cell shows (acc, fts) , where acc is the best accuracy, and fts is the number of significant features used.

n SFs, where the SF ordering is determined by running t-test on each feature using all the objects; (3) **SF-J**: use the first n SFs, as in **SF**, but with the t-test applied only to the training set.

Figure 3 shows the experimental results of FLD-BM and SVM-C10. Clearly, using too few features cannot separate the classes well, while using too many introduces extra noise. The **SF** results show a nearly perfect classification for each classifier in the best case; however, in this case feature selection introduces some bias, as test subjects are included in the selection process. But, it is interesting to see that a feature subset does exist that supports nearly perfect classification. The **SF-J** results always outperform the **PC** for each classifier in terms of the best case. The improvements are about 20% for the left case and 10% for the right case.

The performances of FLD-BM, FLD-1NN, FLD-3NN and FLD-NM are similar, and so are those of SVM-C10 and SVM-C100. But, SVM-C1 underperforms SVM-C10, which indicates the cost of constraint violation needs to be set properly. Table 1 shows the best accuracy using **SF-J** feature selection approach to-




Figure 4: Mean surfaces of (a) left and (b) right groups, both in two different views. Colors indicate magnitude of contribution of each landmark location to the discriminations after PCA and FLD using 14 SFs. Yellow/light color indicates more contribution while green/dark less.

gether with the number of significant features used for each case. SVM-C10 performs the best, with 93% accuracy for the left set and 82% for right, which is competitive with the best result in previous hippocampal studies [3, 5, 11, 12], including our own prior results [11]. Our data set is different from the data sets in previous studies by others, due to a lack of shared data repositories in this domain. However, these are similar results using different techniques on similar types of data.

3.4 Visualization of Group Differences

We introduce a new way to visualize the discrimination between groups. Applying PCA and FLD (possibly with feature selection) as detailed above to a shape set, we get a discriminative value v for each shape \mathbf{x} : $v = \mathbf{x}_\delta^T * B_{pca} * B_{fld} = \mathbf{x}_\delta^T * \mathbf{w}$, where $\mathbf{x}_\delta = \mathbf{x} - \bar{\mathbf{x}}$ is the deformation of \mathbf{x} from the mean shape $\bar{\mathbf{x}}$, B_{pca} consists of a subset of eigenvectors, depending on which PCs are selected, and B_{fld} is the corresponding FLD basis. Thus \mathbf{w} is a column vector that weights the contribution of each deformation element in \mathbf{x}_δ to v . We can map these weights onto the mean surface to show significant discriminative regions. Perfect discriminations are achieved after PCA and FLD using 14 SFs through SF t-tests for both left and right cases; Figure 4 shows the mapping of the weights to surface locations. Mapping results are very similar with adding more features, and show that the head and tail regions tend to have group differences, which matches previous studies.

4 Conclusion

Our proposed classification framework for 3D closed surface objects employs SPHARM expansion and normalization techniques to create a shape descriptor. A PDM is applied to reduce a landmark description to a feature vector, while keeping all the data variance to minimize representation errors. We consider feature selection using original PC orderings as well as SF orderings via a t-test. Several FLD and linear SVM variants are used to perform the classification using selected features. We demonstrate the effectiveness of our framework with a

hippocampal surface classification between controls and schizophrenics. We achieve the best jackknife accuracy of 93%, competitive with the best result in similar studies using different techniques on similar types of data.

SPHARM is a powerful surface modeling approach for arbitrarily shaped but simply connected objects and can deal with protrusions and intrusions. Therefore, our framework can be applied to many other appropriate 3D shape classification problems in computer vision and image processing, such as shape-based diagnosis for other brain structures, organs, or tumors.

References

- [1] Ch. Brechbhlér, G. Gerig, O. Kubler, “Parametrization of closed surfaces for 3D shape description,” *Comp. Vis. and Im. Un.*, 61(2):154–170, 1995.
- [2] T. F. Cootes, C. J. Taylor, D. H. Cooper, J. Haslam, “Active Shape Models—Their Training and Application,” *Comp. Vis. and Im. Un.*, 61(1): 38–59, 1995.
- [3] J. G. Csernansky, S. Joshi, et al., “Hippocampal Morphometry in Schizophrenia by High Dimensional Brain Mapping,” *Proc. Natl. Acad. Sci. USA*, 95:11406–11411, Sep., 1998.
- [4] R. O. Duda, P. E. Hart, D. G. Stork, *Pattern Classification (2nd ed)*, Wiley, New York, NY, 2000.
- [5] G. Gerig, M. Styner, “Shape versus Size: Improved Understanding of the Morphology of Brain Structures,” *MICCAI’2001*, LNCS 2208:24–32, 2001.
- [6] G. Gerig, <http://www.cs.unc.edu/~gerig/pub.html>.
- [7] A. Kelemen, G. Szekely, G. Gerig, “Elastic Model-based Segmentation of 3-D Neuroradiological Data Sets,” *IEEE T. on Med. Im.*, 18(10):828–839, 1999.
- [8] J. Ma, Y. Zhao, S. Ahalt, *OSU SVM Classifier Matlab Toolbox (ver 3.00)*, http://eewww.eng.ohio-state.edu/~maj/osu_svm/, 2002.
- [9] L. Shen, J. Ford, F. Makedon, A. Saykin. Hippocampal Shape Analysis: Surface-based Representation and Classification. *Medical Imaging 2003*, SPIE Proc. 5032:253–264, San Diego, CA, Feb. 2003.
- [10] M. Styner, G. Gerig, S. Pizer, S. Joshi, “Automatic and Robust Computation of 3D Medial Models Incorporating Object Variability”, *Int. J. of Computer Vision*, to appear.
- [11] A. Saykin, L. Flashman, et al., “Principal Components Analysis of Hippocampal Shape in Schizophrenia,” *Int. Con. on Sch. Res.*, 2003.
- [12] S. J. Timoner, P. Golland, R. Kikinis, M. E. Shenton, et al., “Performance issues in shape classification”, *MICCAI’2002*, LNCS 2488:355–362, 2002.