

Statistical Shape Analysis Using Fixed Topology Skeletons: Corpus Callosum Study

Polina Golland¹, W. Eric L. Grimson¹, and Ron Kikinis²

¹ Artificial Intelligence Laboratory, Massachusetts Institute of Technology,
Cambridge, MA 02139, USA
{polina, welg}@ai.mit.edu

² Surgical Planning Laboratory, Brigham and Women's Hospital,
Boston, MA 02115, USA
kikinis@bwh.harvard.edu

Abstract. The goal of this work is to develop an approach to shape representation and classification that will allow us to detect and quantify differences in shape of anatomical structures due to various disorders. We used a robust version of skeletons for feature extraction and linear discriminant analysis (the Fisher linear discriminant and the linear Support Vectors method) for classification. We propose a way to map the classification results back into the image domain, interpreting shape differences as a deformation required to bring a shape from one class to the other. An example of analyzing corpus callosum shape in schizophrenia is reported, as well as the results of the study of the statistical properties of the classifier using cross validation techniques.

1 Introduction

Our goal is to build a framework for statistical shape analysis using classification techniques applied to feature descriptors. We perform shape feature extraction using skeletons. To make the process of skeleton extraction robust to noise and quantization effects of segmentation, we have developed a new variation of the traditional skeletons: *fixed topology skeletons*.

In this paper, we limit ourselves to linear discriminant analysis, comparing performance of two different linear classification methods: the Fisher linear discriminant and the linear Support Vectors methods. Then we present the shape differences between the groups by constructing the shape deformation in the image space that corresponds to the discriminant vector in the feature space.

We tested the approach on corpus callosum data for schizophrenia patients. The results are reported in Sect. 3.

Related Work. Statistical shape modeling combines shape representation with statistical information on how the features vary across population. Principal Component Analysis (PCA) has been used by several authors for capturing statistical properties of the model [3,9]. It was well suited for applications in segmentation and object localization, where the statistical properties of the model

were used to restrict the space of possible deformations of the model. It has also been used in shape analysis [4,8] to reduce the dimensionality of the model and find a decision boundary between the classes. Bookstein [2] used the shape features to align the outlines, but then the features (points along the outline) were analyzed independently of each other. We attempt to use traditional classification methods directly (without going through the dimensionality reduction step) to find the decision boundary.

We use a novel approach to robust skeleton estimation for feature extraction. Skeletons have been introduced in general computer vision several decades ago [1] and have been used extensively for object recognition and localization. In medical image analysis, a scale-space variation of skeletons was introduced and used in various applications by Pizer and colleagues [6].

2 Shape Representation: Fixed Topology Skeletons

Skeletons provide a compact, intuitive representation of a shape that can be used for segmentation, tracking, object recognition, etc. Their major drawback is their high sensitivity to noise in the boundary. There have been proposed many ways to stabilize the skeleton extraction, most of which concentrated on heuristics for pruning the original, noisy skeleton.

For shape analysis of anatomical structures, the general shape of the object is well known ahead of time and the deformations of interest are very small and do not change the global shape of the structure. *Fixed topology skeletons* take advantage of this fact: we fix the structure of the skeleton graph (the skeleton topology) and optimize for the accuracy of the original shape representation over all skeletons of that fixed structure.

Skeleton extraction. For computing the fixed topology skeleton of a shape, we use a distance map, a function that for every point in the image is equal to the distance from the point to the closest point on the boundary of the object. It can be shown that the skeleton is the set of ridge points of the distance map.

We use a snake-like approach for computing the fixed topology skeleton of a shape. The set of skeleton points defines a continuous curve that represents the skeleton. We initialize the snake at the end-points of the traditionally defined skeleton [1,6], and then use the distance map gradient to "drive" the snake. Additional regularization is required to keep the curve smooth. Formally, the update rule is

$$\mathbf{x}^{t+1} = \sigma(\mathbf{x}^t + \nabla D(\mathbf{x}^t)),$$

where \mathbf{x}^t is the set of point coordinates on the curve at time t , ∇D is the gradient of the distance map computed at the locations corresponding to the points of the curve, and σ is the smoothing operator. The curve has to be resampled every few iterations to maintain uniform distribution of the points along the curve. We stop the iterations when the curve starts oscillating around the ridge.

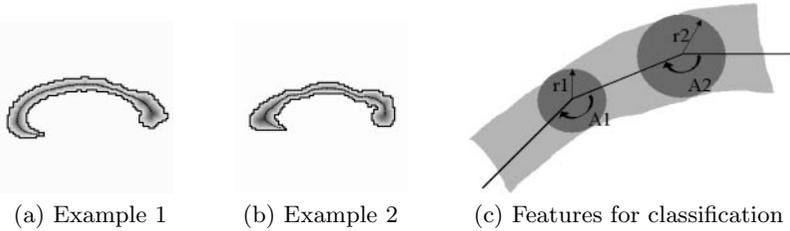


Fig. 1. Skeleton extraction. (a) and (b) show the distance map (darker color corresponds to higher distance from the boundary) and the skeleton extracted for two different cases from the data set; (c) features used for classification: curvature angle and shape width

To find the best skeleton, we estimate skeletons for different initial pairs of points and chose the one that describes the shape the best ([7] contains more details on the algorithm). Figs. 1(a) and (b) show corpus callosum skeletons computed for two different cases in our data set.

Feature extraction. Once the skeleton is computed, we sample the skeleton curve uniformly by arc length and measure two values at every sample point (Fig. 1c): the angle between two adjacent segments in the sampled skeleton and the shape width at the sample point. These two features are invariant under rigid transformations and are therefore well suited for shape description. The number of sampling points on the skeleton determines the level of detail captured by the feature vector.

3 Classification Results

Classification methods. We tested two different linear discriminant techniques on the same data set, namely the Fisher discriminant function [5] and the linear Support Vectors classifier [10]. Given two classes of feature vectors $\{\mathbf{x}\}$, any linear learning method searches for weight vector \mathbf{w} that maximizes ‘spread’ between the *projected* points $x = \mathbf{w}^T \mathbf{x}$. The difference between different linear techniques is in how they define spread, or separation, between the classes.

To find an optimal number of features, we use cross-validation. Since our data set is small, we had to resort to *leave-one-out* cross-validation: one case was left out of the training set and then used as a test set. Repeated for all the cases in the data set, this yields an estimate of the generalization accuracy of the method. We report cross-validation results later in this section.

Data. We tested our approach on corpus callosum images for two groups: schizophrenia patients and normal controls. We used two data sets, combined into one in our experiments (see Acknowledgments for more info). The combined data set contains scans of 30 schizophrenia patients (SZ) and of 36 normal controls (NC). We also performed testing on those data sets separately with results very similar to those obtained with the combined data set.

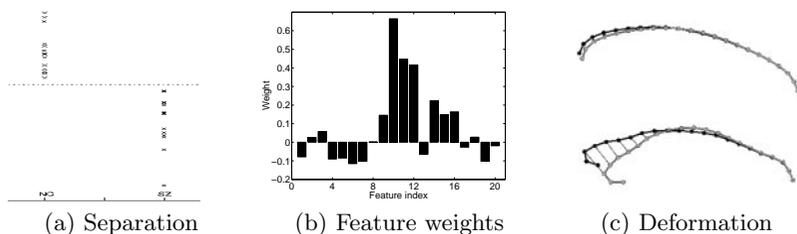


Fig. 2. Classification results based on 20 feature points: (a) separation between the two groups when projected onto \mathbf{w} and (b) weights (components of \mathbf{w}) for the features along the curve, in the posterior-to-anterior order; (c) deformation implied by the discriminant vector, applied to the mean of NC group (top) and to an individual case (bottom). Black corresponds to the original shape, gray indicates the result of the deformation

Classification results. Figure 2(a) shows the results of Support Vectors classification using 20 points along the skeleton. We can see that for this number of features, a perfect separation between the two classes was achieved. Figure 2(b) shows the weights corresponding to the angle features (ordered from posterior to anterior). The weights change smoothly as we move along the skeleton, and most of the weight is concentrated in the middle part of the skeleton. This suggests that the middle ridge is where most of the shape differences take place in this case.

We can also provide a direct interpretation of this result in the image domain. Since projecting onto weight vector \mathbf{w} separates the two classes, negating the component of any feature vector \mathbf{x}_i from the original data set along \mathbf{w} should bring that vector over the threshold into the other class:

$$\begin{aligned} \mathbf{x} &= \mathbf{x}_\perp + (\mathbf{w}^T \mathbf{x})\mathbf{w}, \\ \tilde{\mathbf{x}} &= \mathbf{x}_\perp - (\mathbf{w}^T \mathbf{x})\mathbf{w}. \end{aligned}$$

We can apply this operation to any data point \mathbf{x}_i in one of the classes and then reconstruct the skeleton using the resulting feature vector $\tilde{\mathbf{x}}_i$. Thus linear classification in the feature domain can be mapped into a shape deformation in the image domain.

Figure 2(c) shows the deformation applied to two different skeletons. The first example (top) shows a ‘mean’ normal control skeleton. It was constructed by averaging the features at the 20 points along the skeleton and reconstructing a skeleton from the resulting feature vector. The second example (bottom) shows a skeleton for one of the normal control subjects with the deformation implied by the classifier. We can see that the corpus callosum shape is more ‘bent’ for schizophrenia group. In other words, we would have to bend the normal corpus callosum further to make it look more like corpus callosum of a schizophrenia patient.

Cross-validation. Figure 3(a) shows learning accuracy, that is the classification accuracy, when the test set was the same as the training set. We can see that

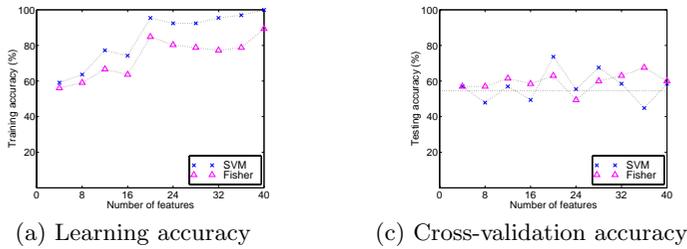


Fig. 3. Cross-validation results

the Support Vectors method outperforms the Fisher linear discriminant, which we believe is because it makes fewer assumptions on the underlying distributions of the classes. As the number of feature points used for classification grows, the data becomes more separable and the accuracy improves.

Cross-validation is used to find the optimal number of feature points for be used for shape description of corpus callosum, as well as to test the generalization power of the classifier. Figure 3(b) shows the classification accuracy for leave-one-out cross-validation experiment. The dotted line shows the ‘baseline’, or the classification accuracy one would get by guessing. We can see that both methods achieve better than guessing accuracy. The best accuracy was achieved by Support Vectors method for 20 feature points. Thus that was reported as the best number of points.

The classification accuracy for cross-validation is significantly lower than for learning. There are several reasons for that. As the number of feature points grows, the data becomes more sparse in the feature space, and thus it is easier to separate between the classes, but we get poor generalization, as new examples fall into previously empty regions of the feature space. Another reason for lower testing accuracy could be that the classes are not truly separable¹.

Another question that should be addressed is the number of features. It seems that the optimal number of features is comparable with the number of cases in the data set. But it does not mean that we are fitting a model with that many *independent* parameters to the data. In fact, the features highly correlate with their neighbors along the skeleton. Another point to confirm this is the fact that adjacent points on the skeleton get similar weights (Fig. 2b).

4 Conclusions & Acknowledgments

We presented an approach for shape based classification of anatomical structures. It uses statistical learning techniques for investigating the differences between two groups of examples of the same anatomical structure. In this work,

¹ Implying that one could not provide a reliable diagnosis of schizophrenia based on the shape of corpus callosum alone, but only about 70% accurate estimate. But combined with analysis of other structures, it might provide a significant improvement in detecting and quantifying shape pathologies in the brain of schizophrenia patients.

we limited ourselves to using linear classifiers. We tested two different linear classification techniques: the Fisher linear discriminant and the linear Support Vectors classification.

The shape representation is also a crucial component of the system. It maps the images into points in the feature space in which the classification is performed, and also provides an interpretation of the classification results in terms of the shape deformation. We use skeletons for extracting the shape features. They provide a robust, intuitive representation of the shape, and are capable of capturing shape variations between the groups reported in the paper.

Based on the experimental results, we conclude that the shape of corpus callosum is different in schizophrenia with higher curvature of the shape. The cross-validation provided the optimal number of the feature points, as well as an estimate of the classification accuracy on the new examples.

Acknowledgments. The authors would like to thank M. Frumin and M. E. Shenton of Harvard Medical School and Brigham and Women's Hospital for providing DataSet1 and G. Gerig of University of North Carolina for providing DataSet2.

This work was supported by Mitsubishi Electric Research Laboratories and NSF IIS-9610249 grant. The project of acquisition and segmentation images for DataSet1 at Brigham and Women's Hospital was supported by Veterans Administration Psychiatric Research/Neuroscience Fellowship, Dupon Warren Fellowship from the Consolidated Department of Psychiatry, NIH grants MH 50740 and MH 01110.

References

1. Blum, H.: Biological shape and visual science. *Journal of Theoretical Biology* **38** (1973) 205-287
2. Bookstein, F.L.: Landmark methods for forms without landmarks: morphometrics of group differences in outline shape, *Medical Image Analysis* **1** (1996) 225-243
3. Cootes, T.F., and Taylor, C.J. Active Shape Models - 'Smart Snakes', In Proc. British Machine Vision Conference, Springer-Verlag (1992) 266-275
4. Csernansky, J.G. *et al*, Hippocampal morphometry in schizophrenia by high dimensional brain mapping. *Proc. Nat. Acad. of Science* **95** (1998) 11406-11411
5. Duda, R.O. and Hart, P.E.: *Pattern Classification and Scene Analysis*, John Wiley & Sons, 1973
6. Fritsch, D.S. *et al*, Stimulated Cores and their Applications in Medical Imaging. In Proc. IPMI'95 (1995) 365-368
7. Golland, P., and Grimson, W.E.L.: Fixed topology skeletons, AI Lab Memo, Massachusetts Institute of Technology, Submitted to ICCV'99, 1999
8. Martin, J., Pentland, A., and Kikinis, R.: Shape Analysis of Brain Structures Using Physical and Experimental Models, In Proc. CVPR'94 (1994) 752-755
9. Székely, G. *et al*, Segmentation of 2D and 3D objects from MRI volume data using constrained elastic deformations of flexible Fourier contour and surface models. *Medical Image Analysis* **1** (1996) 19-34
10. Vapnik, V.N.: *Statistical Learning Theory*, John Wiley & Sons, 1998