ISSN 2007-9737

Hierarchical Contour Shape Analysis

Daniel Valdes-Amaro[1], Abhir Bhalerao[2]

[1] Benemerita Universidad Autonoma de Puebla, Faculty of Computer Science, Puebla, Mexico

[2] University of Warwick, Department of Computer Science, Coventry, UK

daniel.valdes@cs.buap.mx, abhir.bhalerao@dcs.warwick.ac.uk

Abstract. This paper introduces a novel shape representation which performs shape analysis in a hierarchical fashion using Gaussian and Laplacian pyramids. A background on hierarchical shape analysis is given along with a detailed explanation of the hierarchical method, and results are shown on natural contours. A comparison is performed between the new method and our proposed approach using Point Distribution Models with different shape sets. The paper concludes with a discussion and proposes ideas on how the new approach may be extended.

Keywords. Shape analysis, shape representation, Gaussian pyramids, shape models, brain contours.

1 Introduction

Hierarchical Shape Analysis can be regarded as a technique able to derive and quantify correlated behavior between any number of structures [14]. Among these techniques, it is possible to find the Hierarchical ASM [8], where the premise of the method is that a small number of training samples can be used to estimate the covariation of patterns of multiple variables, allowing their hierarchical model to capture global (coarse) and local (fine) shape details. This is done by representing, in a hierarchical fashion, the shapes in terms of a wavelet transform followed by a Principal Component Analysis (PCA) of the computed coefficients.

In [17] a technique called Partitioned Active Shape Model (PASM) is introduced, where these 3D PASMs can be regarded as partitioned representations of 3D ASMs. Here a mesh is partitioned into 'tiles' or surface patches, then PCA is applied to the coordinates of the tile vertices. Next, training samples are projected as curves in a single hyperspace so that the deformed points are fitted into an allowable region of the model using a curve alignment scheme.

Rao *et al.* [14] presented an approach where two well-known multivariate statistical techniques are used to investigate the statistical variation of brain structures. Canonical Correlation Analysis (CCA) is used to quantify the correlations between the number of different brain structures, and Partial Least Squares Regression (PLSR) is performed over shapes of different structures to predict unknown shapes from unseen subjects, given known shapes of other structures from that subject. Additionally, CCA and PLSR facilitate the embedding of statistical shape modeling of the brain within a hierarchical framework, since they can be used to extract and quantify correlated behavior between any number of brain substructures at multiple scales.

More recently, Yu *et al.* [16] introduced a method that extracts shape features and conducts statistical analysis using a procedure that registers and normalizes cortical surfaces, as well as decompose them using spherical wavelets. The wavelet coefficients obtained are used as shape features to study the folding pattern at different spatial scales and locations, as the underlying wavelet basis functions have local support in the space and frequency domains. Then, the patterns of cortical shape variation are studied using PCA allowing the correlation of these shape variations with age and neuropsychological measurements at different spatial scales.

In Table 1 an overview of these methods is provided. Each entry in the table shows information about the modeled object of interest, the image modality, and the techniques used for the work.

Table 1. Overview of Shape Analysis methods and tasks solved using Hierarchical Shape Models. Key: PCA = Principal Component Analysis; PDM = Point Distribution Model; PLSR = Partial Least Squares Regression; CCA = Canonical Correlation Analysis; SPHARM = Spheral Harmonics

Authors	Object(s) of interest	Core technique(s)
Davatzikos *et al.* (2003)	Corpus Callosum	Wavelet transform PCA
Zhao *et al.* (2005)	Lateral ventricle Left Thalamus Left Hippocampus	PDM PCA Curve Alignment model fitting
Rao *et al.* (2006)	Lateral ventricle Pallidum Caudate Putamen Thalamus Amygdala Hippocampus Accumbens Brain Stem	 PLSR CCA
Yu *et al.* (2007)	White Matter Grey Matter	SPHARM PCA

In this paper a new shape model that derives shape information in a hierarchical fashion is presented. To create a multilevel analysis, the method employs the same ideas as in Burt and Adelson's method for images [3], i.e. that of using an iterative process to encode the signal to generate a pyramid data structure, which is equivalent to sampling an image with Laplacian operators of many sizes, hence the name Laplacian pyramid. Analogously, in this work, the idea is that the Laplacian pyramid encodes *shape* variation, so each level encodes the variation between successively smoothed versions of the input shape. We explain the basic principles behind the method and use examples to illustrate its construction. Brain and leaf data sets are then used to produce results of applying the new model. The aim is to show that new hierarchical representation is more compact than a PCA model. Our experiments confirm this hypothesis, which indicates that the method has utility when few training samples are available to model complicated contours. In particular, the method is well suited to natural shapes consisting of repeating or self-similar patterns of variation.

2 Generating Hierarchical Shape Models

In the early 1980s, Burt and Adelson [3] introduced an image encoding technique that uses local operators at many scales as basis functions. The iterative process to encode the image generates a pyramid data structure, which is equivalent to sampling an image with Laplacian operators of many sizes, hence the name Laplacian pyramid. In this section a novel method for shape analysis is introduced based on the idea of Laplacian operators for image enhancement and image coding.

2.1 The Contour Laplacian Pyramid as a Compact Shape Code

As in Burt and Adelson's method [3], the approach applies a Gaussian pyramid $\mathcal{GP}$ as a first step. To generate a $\mathcal{GP}$ for any shape

$$\mathcal{S} = (x_1, y_1, x_2, y_2, \ldots, x_n, y_n)^T, \qquad (1)$$

it is 'reduced' using a Gaussian function that acts as a low pass smoothing filter and reduces the number of points. Such an operation is performed by the function $REDUCE$ which can be defined as

$$REDUCE(\mathcal{S}) = \Downarrow(X(i,\sigma), Y(i,\sigma)). \qquad (2)$$

Here, $X(i,\sigma) = x(i) \circledast g(i,\sigma)$ and $Y(i,\sigma) = y(i) \circledast g(i,\sigma)$, where $\circledast$ denotes the convolution operator and $g(i,\sigma)$ is a Gaussian of width σ [13]. The operator $\Downarrow$ denotes downsampling of the shape by a factor of two. Hence,

$$\mathcal{GP}(i) = REDUCE(\mathcal{S}), \qquad (3)$$

where $i = 0, \ldots, l$ and l is the number of levels of the pyramid. Fig. 1 presents an example of a Gaussian pyramid for a leaf and a brain white matter contour. As the $REDUCE$ operator is repeatedly applied, the shape is successively blurred and reduced in size until a small circle results at a higher pyramid level above the starting level.

The next step is the construction of the Laplacian pyramid $\mathcal{LP}$, where each level is the difference between two adjacent levels of the Gaussian pyramid. First, the function $EXPAND$ is defined as:

$$EXPAND(\mathcal{S}) = \Uparrow(\mathcal{GP}(i)) \circledast g(i,\sigma), \qquad (4)$$

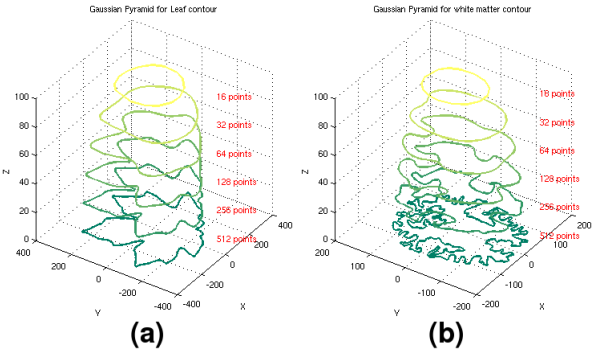

Fig. 1. Example of Gaussian pyramid for **(a)** leaf and **(b)** brain white matter contour

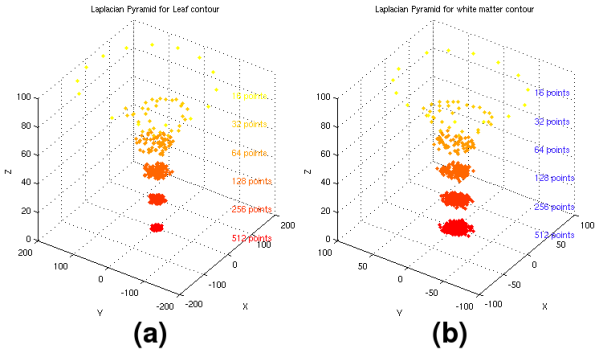

Fig. 2. Example of Laplacian pyramid for **(a)** leaf and **(b)** brain white matter contour

where $\Uparrow$ denotes an operator that up-samples the shape by a factor of two. Then, each level of the Laplacian pyramid $\mathcal{LP}$ is given by

$$\mathcal{LP}(i) = \mathcal{GP}(i) - EXPAND(\mathcal{GP}(i+1)). \quad (5)$$

An $EXPAND$ operator is generally implemented by nearest neighbor up-sampling of the input followed by a smoothing operation. The same smoothing kernel used for the $REDUCE$ operation is reapplied to the up-sampled signal.

Since there is no level $i+1$ of $\mathcal{GP}$ to serve as a prediction level for $\mathcal{GP}(l)$, we say

$$\mathcal{LP}(l) = \mathcal{GP}(l). \quad (6)$$

Fig. 2 presents an example of a Laplacian pyramid for a leaf and a brain white matter contour. The output is harder to interpret as the 'contour' is no-longer the original shape at any level, but represents the difference between adjacent Gaussian contours, i.e. each Laplacian pyramid contour is the *detail*. Note this is analogous to a wavelet signal analysis consisting of a low-frequency approximation and a high-frequency detail.

It is then possible to recover the original contour by expanding $\mathcal{LP}(l)$ once and adding it to $\mathcal{LP}(l-1)$, then expanding again and adding it to $\mathcal{LP}(l-2)$, and so on until level 0 (the bottom of the pyramid) is reached and $\mathcal{GP}(0)$ is recovered (the level that corresponds to the original contour). This procedure simply reverses the steps in the Laplacian pyramid generation. From Equation 5 we observe that

$$\mathcal{GP}(i) = \mathcal{LP}(i) - EXPAND(\mathcal{GP}(i+1)). \quad (7)$$

Fig. 3 presents an example of the reconstruction of the Gaussian pyramid for a leaf and a brain white matter contour, respectively. Fig. 4 shows an

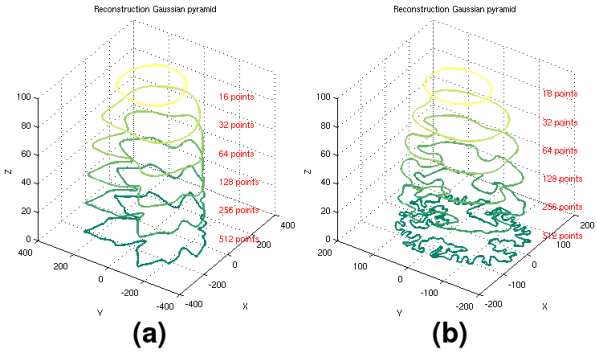

Fig. 3. Example of the reconstruction of the Gaussian pyramid for **(a)** leaf and **(b)** brain white matter contour

example of the construction of the Gaussian and Laplacian pyramids with a contour $\mathcal{S}$ of 512 points. Unsurprisingly, the reconstructions using the top-level of the GP and successive contour differences from the LP result in the reconstructions of the GP, and the figures look identical to the GP figure.

Being able to decompose a shape into approximation and detail in this manner has a number

ISSN 2007-9737

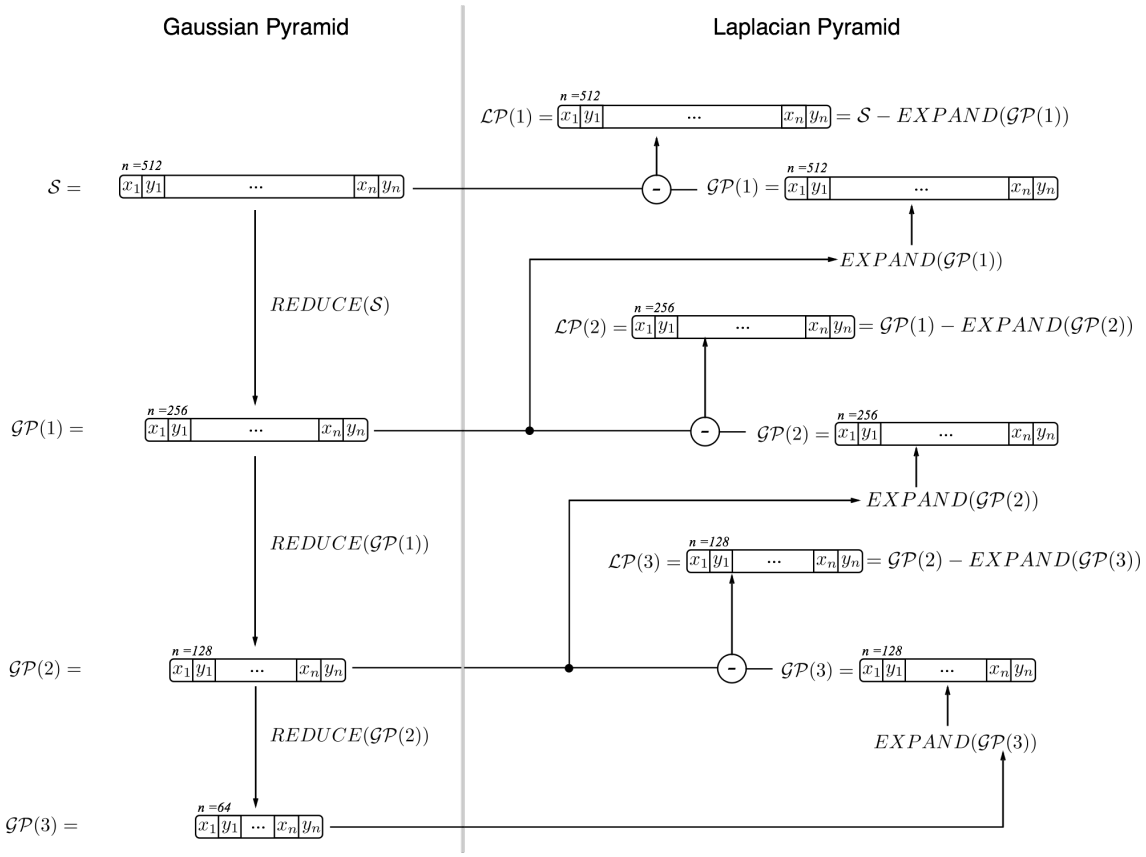

Fig. 4. Example of the construction of the Gaussian and Laplacian pyramids with a contour $\mathcal{S}$ (Equation 1) of 512 points. On the left side the Gaussian pyramid ($\mathcal{GP}$) construction process is depicted, where the shape is successively blurred and reduced in size using the $REDUCE$ operator, until a small circle results at a higher pyramid level above the starting level. Next, on the right side the construction of the Laplacian pyramid ($\mathcal{LP}$) is shown. In this step, the $EXPAND$ operator is used to create the different levels, but despite it is no longer the original shape at any level, each Laplacian pyramid contour represents the 'detail' of the shape

of useful properties. As with other wavelet signal representations, they can be used for denoising and compression using simple threshold methods (shrinkage) since the analysis process acts to decorrelate the input signal.

For example, if the shape contour were to contain noise, we could threshold out Laplacian pyramid shape coefficients below some absolute value and this prior to reconstruction, which would remove noise. In this work, we focus on using Principal Component Analysis (PCA) to compactly represent the shape by encoding shapes over the multiple Laplacian pyramid levels.

2.2 Deriving Shape Information from Laplacian Pyramids

Intuitively, the idea is that the shape variation is encoded in the Laplacian pyramid, so each level encodes the shape variation among the different levels of the Gaussian pyramid, i.e. at different resolutions of detail.

Given a set of shapes $\Phi = \mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_n$, the first step, as in any other shape model, is to align them. This is necessary to remove *pose* variation, this is known as Procrustes Analysis. In this work Generalized orthogonal Procrustes analysis (GPA) [11] is used, where k sets can be aligned to one

target shape or aligned to each other. So, in general for GPA, the process is initiated by selecting the first shape in the set to be the approximate mean shape. Then, a new approximate mean is calculated from the aligned shapes, followed by a test of convergence based on the residual error with the mean.

Once the shape set has been aligned, we proceed to model the shape information by covariance analysis. Each level of the Laplacian pyramid $\mathcal{LP}$ is collapsed to a stacked vector of x and y coordinates, and the different levels are concatenated forming $\mathcal{LV}$:

$$\mathcal{LV} = [\mathcal{GP}(l), \mathcal{L}_{(l-1)}, \mathcal{L}_{(l-2)}, \ldots, \mathcal{L}_{0)}]. \quad (8)$$

where $\mathcal{L}$ represents a level from $\mathcal{LP}$.

Then, a covariance matrix $\mathbf{C}$ is created by the outer product of each $\mathcal{LV}$:

$$\mathbf{C} = \frac{1}{n}\sum_{i=1}^{n}[\mathcal{LV}_i][\mathcal{LV}_i]^T. \quad (9)$$

Then, using equation

$$\mathbf{C}\psi = \lambda\psi \quad (10)$$

over $\mathbf{C}$, we obtain the corresponding eigenvectors and eigenvalues by PCA. The eigenvalues or *modes of variation* effectively capture the variability of the set. Larger eigenvalues are associated with the principal modes which model the principal variation. Detail and noise are represented by the minor modes associated with the smallest eigenvalues [10]. Additionally, the modes of variation can be used in an alternative way to evaluate their importance. The magnitudes of the principal non-zero eigenvalues can be plotted against the number of principal mode, sorted by decreasing eigenvalue. In this type of *eigenshape* analysis, the plot can be used to compare the performance of the methods in the following way: from the eigenvalues plots for different data sets or methods, we observe that the more compact is the shape space, the faster the magnitudes of the eigenvalues tend towards zero [2].

2.3 Point Distribution Models

A *Point Distribution Model* (or PDM) is a way to represent shapes through a model that reflects the position of labeled points [7]. The model also represents the mean geometry of a shape and derives statistical modes of geometric variation inferred from a training set of shapes. It has become a standard in computer vision for statistical shape analysis and especially for segmentation purposes on medical images and has led to the creation of two important techniques: Active Shape Models (ASM) [6] and Active Appearance Models (AAM) [5]. A comparison between both models can be found in [4].

The method can be generalized as follows. First we need to obtain a training set of outlines with enough landmarks so they are able to sufficiently approximate the geometry of the original shapes (Fig. 5-a). Next, an alignment of the landmark sets using GPA is performed (Fig. 5-b). The idea behind this is that the shape information is not related to affine pose parameters, so they need to be removed. Having this, a mean shape can now be computed by averaging the aligned landmark positions (Fig. 5-b). PCA computes the eigenvectors and eigenvalues of the training set using a covariance matrix. Each eigenvector describes a principal mode of variation along the set, the corresponding eigenvalue indicating the importance of this mode in the shape space scattering (Fig. 5-c). If correlation is found between landmarks, then the total variation of the space is concentrated on the first few eigenvectors which present a very rapid decrease in their corresponding eigenvalues. Otherwise, if no correlation is found, that might suggest that the training set has no variation or that the pose of the landmarks has not been properly removed. Finally, using the set of generated eigenvectors and eigenvalues, any shape of the training set can be approximated using the mean shape and a weighted sum of deviations obtained from the modes (Fig. 5-d).

The generated eigenvectors can be seen as a sequence of vectors associated to the corresponding landmarks, where each comprises a mixture of shape variation for the whole shape. Here, the model consists of the mean positions of these

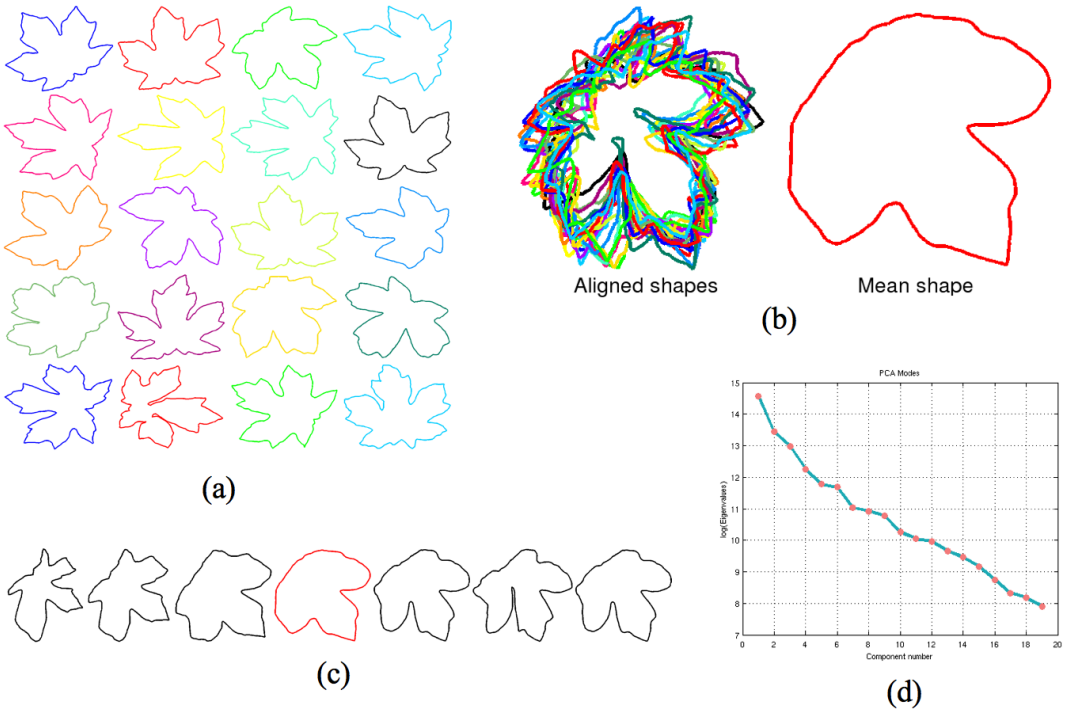

Fig. 5. Point Distribution Model: (a) training set of contours, (b) aligned shapes after GPS and set mean shape, (c) from the mean shape it is possible to reconstruct any shape of the set adding the proper modes of variation, (d) log-plot of the eigenvalues against the number of modes

points and the main modes of variation, which describe how the points tend to move from the mean. The most important idea behind the PDMs is that eigenvectors can be linearly combined to create new shape instances that will be similar to any in the training set [7].

Fig. 6. Leaf types used in the experiments: (a) Macrophyllum and (b) Kelloggii

3 Comparative Assessment: Laplacian Hierarchical Shape Model and PDM

In this section we present experiments that compare the novel Hierarchical Shape Model (HSM) and the Point Distribution Model. For both experiments, different levels of the Gaussian pyramid were used, and the idea was to compare the compactness of the eigenvalues from both models. In each of the following figures the resulting eigenmodes for the HSM are plotted in shades of green and for the PDM in red.

The first data set contains 50 leaves of the type Macrophyllum (Fig. 6-a) as well from [9]. Fig. 7 corresponds to the HSM using 3 levels of the Gaussian pyramid, Fig. 8 to 4 levels, Fig. 9 to 5, and finally Fig. 10 to 6 levels.

The second data set contains 60 brain white matter contours from simulated digital brain phantom images, and each digital brain was created

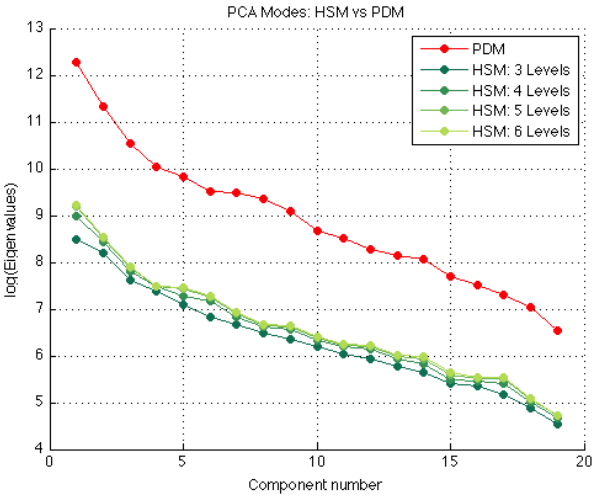

Fig. 7. Macrophyllum leaves set of 20 shapes: 4 different levels of the Gaussian pyramid in shades of green, and in red the plot for the PDM. Plot of eigenmodes against the number of principal modes

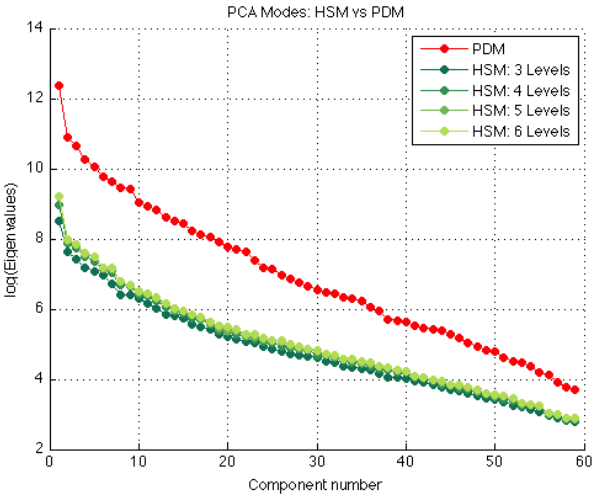

Fig. 9. Macrophyllum leaves set of 60 shapes: 4 different levels of the Gaussian pyramid in shades of green, and in red the plot for the PDM. Plot of eigenmodes against the number of principal modes

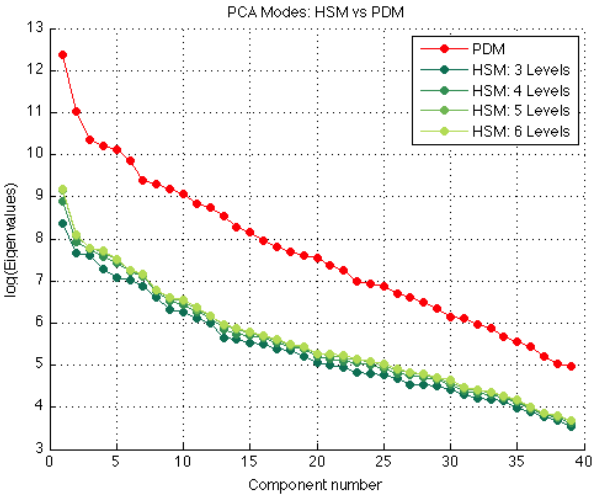

Fig. 8. Macrophyllum leaves set of 40 shapes: 4 different levels of the Gaussian pyramid in shades of green, and in red the plot for the PDM. Plot of eigenmodes against the number of principal modes

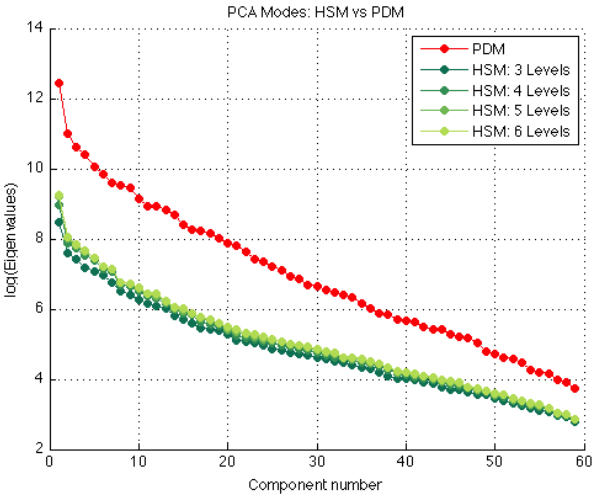

Fig. 10. Macrophyllum leaves set of 80 shapes: 4 different levels of the Gaussian pyramid in shades of green, and in red the plot for the PDM. Plot of eigenmodes against the number of principal modes

by registering and averaging four T1, T2, and PD-weighted MRI scans from normal adults [1].

Again as in the previous data set, each Figure (11, 12, 13 and 14) corresponds to the HSM using 3, 4, 5, and 6 levels of the Gaussian pyramid, respectively.

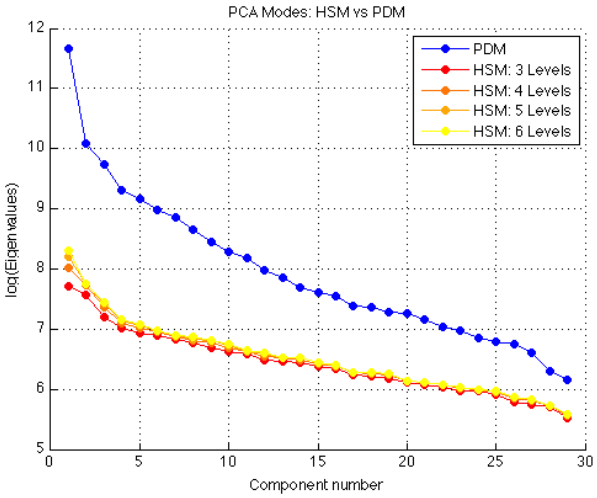

Fig. 11. Brain white matter set of 30 shapes: 4 different levels of the Gaussian pyramid in warm colors, and in blue the plot for the PDM. Plot of eigenmodes against the number of principal modes

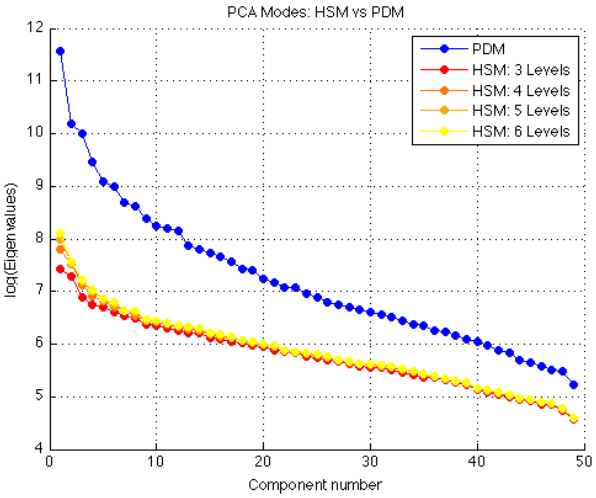

Fig. 13. Brain white matter set of 50 shapes: 4 different levels of the Gaussian pyramid in warm colors, and in blue the plot for the PDM. Plot of eigenmodes against the number of principal modes

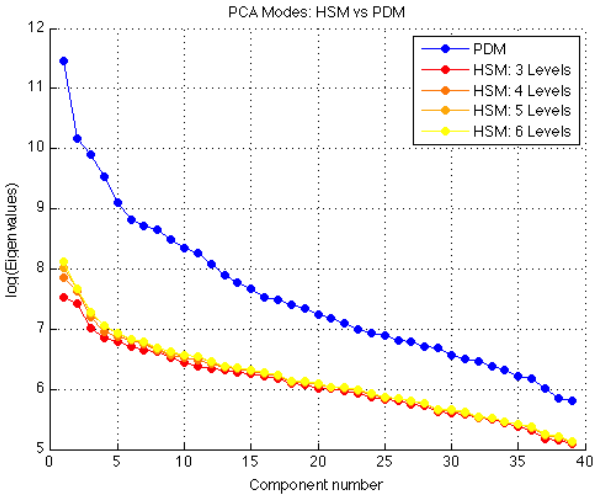

Fig. 12. Brain white matter set of 40 shapes: 4 different levels of the Gaussian pyramid in warm colors, and in blue the plot for the PDM. Plot of eigenmodes against the number of principal modes

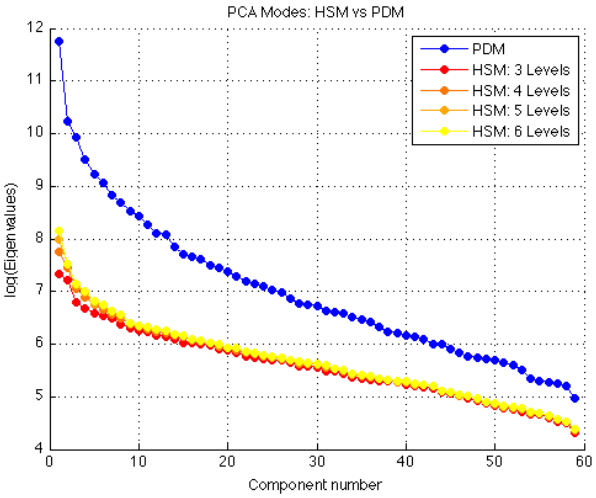

Fig. 14. Brain white matter set of 60 shapes: 4 different levels of the Gaussian pyramid in warm colors, and in blue the plot for the PDM. Plot of eigenmodes against the number of principal modes

4 Evaluation and Discussion

In this paper a new shape model that derives shape information in a hierarchical fashion was presented. Hierarchical Shape Analysis can be regarded as a technique able to derive and quantify correlated behavior among any number of structures. The resulting plots show that even using dif-

ferent levels of the Laplacian pyramid and different sizes of the training sets, the plots do not present significative variation in their compactness. Hence, it is possible to conclude that the compactness of the method is consistent, and therefore using few levels or a small set of shapes will not affect the performance.

Next, we compared its compactness against the PDM results from leaf contours (Figures 7 to 10) and from white matter shapes (Figures 11 to 14). The results illustrate that, in most cases, the HSM is more compact than the PDM since it is possible to find significative variation in the compactness of the eigenmode plots. Also, the accuracy in deriving shape information outperforms the proposed PDM in some cases. This can be graphically assessed in the resulting plots, where none of the plots using different levels of the HSM were less compact than the ones produced by the PDM. Moreover, it can be seen that even using a few tree levels of the HSM good compactness can be achieved. This means that the HSMs can be at least as good and compact as the PDM, and consequently is worthy for performing efficient shape analysis but in a different way.

Currently there remains a scope for extending the presented technique in two ways. The first is to adjust it to local shape analysis by taking parts of the contour at different scales of smoothing according, for example, a Curvature Scale Space (CSS) description (see [15]). The other direction will be to extend it to three dimensions using surface patches. For example, in [12] it is shown that by convolving local parameterizations of the surface with 2D Gaussian filters iteratively it is possible to obtain smoothed versions of the patches. We believe this should be sufficient to derive the Gaussian and Laplacian pyramids for a surface HSM.

Acknowledgment

D. Valdes-Amaro would like to thank SEP-PROMEP (PROMEP/103.5/12/8136) for the financial support given to this research.

References

1. **Aubert-Broche, B., Griffin, M., Pike, G. B., Evans, A. C., & Collins, D. L.** (**2006**). Twenty new digital brain phantoms for creation of validation image data bases. *IEEE Transactions on Medical Imaging*, Vol. 25, pp. 1410–14163.
2. **Bhalerao, A. & Wilson, R.** (**2005**). Local Shape Modelling Using Warplets. **Kälviäinen, H., Parkkinen, J., & Kaarna, A.**, editors, *Image Analysis, 14th Scandinavian Conference, SCIA 2005, Joensuu, Finland, June 19-22, 2005, Proceedings*, volume 3540 of *Lecture Notes in Computer Science*, Springer, pp. 439–448.
3. **Burt, P. J. & Adelson, E. H.** (**1983**). The laplacian pyramid as a compact image code. *IEEE Transactions on Communications*, Vol. COM-31, No. 4, pp. 532–540.
4. **Cootes, T. F., Edwards, G., & Taylor, C.** (**1999**). Comparing Active Shape Models with Active Appearance Models. *Proceedings of the British Machine Vision Conference, BMVC 1999, University of Nottingham, September 13-16, 1999.*, BMVA Press, pp. 173–182.
5. **Cootes, T. F., Edwards, G. J., & Taylor, C. J.** (**1998**). Active Appearance Models. *5th European Conference on Computer Vision*, volume 1407, Springer, Berlin, pp. 484–498.
6. **Cootes, T. F. & Taylor, C. J.** (**1992**). Active Shape Models: Smart Snakes. *British Machine Vision Conference*, pp. 267–275.
7. **Cootes, T. F., Taylor, C. J., Cooper, D. H., & Graham, J.** (**1992**). Training models of shape from sets of examples. *Proc. British Machine Vision Conference*, Springer, Berlin, pp. 266–275.
8. **Davatzikos, C., Tao, X., & Shen, D.** (**2003**). Hierarchical active shape models, using the wavelet transform. *IEEE Transactions on Medical Imaging*, Vol. 22, No. 3, pp. 414–423.
9. **Dietterich, T. G.** (**2002**). Isolated leaves dataset, oregon state university web resource, url: http://web.engr.oregonstate.edu/ tgd/leaves/.
10. **Fukunaga, K. & Koontz, W. L. G.** (**1970**). Applications of the Karhunen-Loeve expansion to feature selection and ordering. *IEEE Transactions on Computers*, Vol. C-19, pp. 311–318.
11. **Gower, J. C.** (**1975**). Generalized Procrustes Analysis. *Psychometrika*, Vol. 40, pp. 33–51.

12. **Mokhtarian, F., Khalili, N., & Yuen, P.** (**2002**). Estimation of error in curvature computation on multiscale free-form surfaces. *International Journal of Computer Vision*, Vol. 48, No. 2, pp. 131–149.
13. **Mokhtarian, F. & Mackworth, A. K.** (**1995**). A theory of multiscale, curvature-based shape representation for planar curves. *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 14, No. 8, pp. 789–805.
14. **Rao, A., Aljabar, P., & Rueckert, D.** (**2008**). Hierarchical statistical shape analysis and prediction of sub-cortical brain structures. *Medical Image Analysis*, Vol. 12, pp. 55–68.
15. **Valdes-Amaro, D. A. & Bhalerao, A.** (**2008**). Local Shape Modelling for Brain Morphometry using Curvature Scale Space. **McKenna, S. & Hoey, J.**, editors, *Proceedings of the 12th Annual Conference on Medical Image Understanding and Analysis 2008*, British Machine Vision Association, pp. 64–68.
16. **Yu, P., Grant, P. E., Qi, Y., Han, X., Ségonne, F., Pienaar, R., Busa, E., Pacheco, J., Makris, N., Buckner, R. L., Golland, P., & Fischl, B.** (**2007**). Cortical Surface Shape Analysis Based on Spherical Wavelets. *IEEE Trans. Medical Imaging*, Vol. 26, No. 4, pp. 582–597.
17. **Zhao, Z., Aylward, S. R., & Teoh, E. K.** (**2005**). A novel 3D Partitioned Active Shape Model for Segmentation of Brain MR Images. **Duncan, J. S. & Gerig, G.**, editors, *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2005, 8th International Conference, Palm Springs, CA, USA, October 26-29, 2005, Proceedings, Part I*, volume 3749 of *Lecture Notes in Computer Science*, Springer, pp. 221–228.

Daniel Valdes-Amaro is part time professor in Computer Science at the Autonomous University of Puebla, Mexico. He got a BCS from the Faculty of Computer Science at the Autonomous University of Puebla, and later a PhD in Computer Science from the Computer Science Department at the University of Warwick in Coventry, UK in 2010. His current research interests are in Statistical Shape Analysis particularly Local Shape Modelling, Fractal Geometry applied to Computer Vision, Spectral Clustering Analysis and Computer Graphics oriented to Augmented Reality and Multimedia.

Abhir Bhalerao is an Associate Professor (Senior Lecturer) in Computer Science. He joined as staff in 1998 having completed 5 years as a post-doctoral research scientist with the NHS and Kings Medical School, London including two years as a Research Fellow at Harvard Medical School Boston. At Warwick, Dr Bhalerao has been principal and co-investigator on two EPSRC sponsored projects related to multi-resolution segmentation and modelling of vasculature from MR (GR/M82899/01(P)) and use of stochastic methods for statistical image modelling (GR/M75785/01, a project undertaken jointly with Department of Statistics). His current interests are in modelling flow in 3D X-ray imagery, characterisation of the morphology of folding structures in MRI, real-time visualization of tensor data and light-estimation methods from multi-camera scenes. He is the co-founder and was the Research Director of Warwick Warp Ltd., a company specialising in biometric technologies.

Article received on 31/01/2014; accepted on 17/04/2015.
Corresponding author is Daniel Valdes-Amaro.