

ICA vs. PCA Active Appearance Models: Application to Cardiac MR Segmentation

M. Üzümcü¹, A.F. Frangi², M. Sonka³, J.H.C. Reiber¹, B.P.F. Lelieveldt¹

¹Div. of Image Processing, Dept. of Radiology

Leiden University Medical Center, Leiden, The Netherlands

²Division of Biomedical Engineering, Aragon Institute of Engineering Research
University of Zaragoza, Zaragoza, Spain

³Dept. of Electrical and Computer Engineering, the University of Iowa, Iowa, USA
M.Uzumcu@lumc.nl

Abstract. Statistical shape models generally use Principal Component Analysis (PCA) to describe the main directions of shape variation in a training set of example shapes. However, PCA has the restriction that the input data must be drawn from a Gaussian distribution and is only able to describe global shape variations. In this paper we evaluate the use of an alternative shape decomposition, Independent Component Analysis (ICA), for two reasons. ICA does not require a Gaussian distribution of the input data and is able to describe localized shape variations. With ICA however, the resulting vectors are not ordered, therefore a method for ordering the Independent Components is presented in this paper. To evaluate ICA-based Active Appearance Models (AAMs), 10 leave-15-out models were trained on a set of 150 short-axis cardiac MR Images with PCA-based as well as ICA-based AAMs. The median values for the average and maximal point-to-point distances between the expert drawn and automatically segmented contours for the PCA-based AAM were 2.95 and 8.39 pixels. For the ICA-based AAM these distances were 1.86 and 5.01 pixels respectively. From this, we conclude that the use of ICA results in a substantial improvement in border localization accuracy over a PCA-based model.

1 Introduction

Statistical shape models such as Active Shape Models (ASM) [1] or Active Appearance Models (AAM) [2] have shown great potential in image recognition and segmentation tasks [3,4,5]. Such models generally use Principal Component Analysis to describe the main directions of shape variation in a training set of example shapes. PCA is a method that fits a Gaussian hyper-ellipsoid to the data and therefore assumes that the data is normally distributed, which frequently is not the case. This may lead to a wrong statistical description of the dataset and to the generation of implausible shapes. Furthermore, PCA yields eigenvectors that describe global variations: when a parameter corresponding to one vector is changed, the entire shape deforms, which makes it difficult to achieve a good local segmentation.

In this paper, we investigate the use of an alternative decomposition technique, Independent Component Analysis, which does not require a Gaussian distribution of the

input data and yields vectors that describe localized shape variations, and therefore potentially resolves these problems.

To enable the use of ICA in statistical shape modeling, two problems had to be addressed:

- Several different methods exist for calculating independent components, which all give highly similar, yet not identical components.
- In classical PCA, the eigenvectors are ordered according to the associated eigenvalues (variances). In Independent Component Analysis however, the resulting vectors are not ordered.

In previous work, we have evaluated different methods for computing and ordering the Independent Components (ICs). From that work, the methods that perform best were selected, and the main contribution of this paper is the evaluation of Independent Component Analysis for image segmentation using Active Appearance Models. In addition, a quantitative comparison between conventional PCA-based and our ICA approach is presented.

The paper is further organized as follows. In Section 2, Active Appearance Models are briefly explained. In Section 3, Independent Component Analysis is introduced, as well as the ordering method that was developed to sort the Independent Components. In Section 4, AAM segmentation results obtained with PCA and with ICA are presented. Section 5 concludes with a discussion.

2 Active Appearance Models

Statistical Models for image processing, such as the Active Shape Model and the Active Appearance Model were introduced by Cootes et al. [1,2], and can be applied to image segmentation in various domains [3,4,5]. These models describe statistical variations as seen in a set of example images, in which corresponding landmark points are annotated. The shapes, which are spanned by the landmarks, are aligned using Procrustes analysis to compensate for translation, rotation and scaling differences between the shape samples. The mean shape of this set of aligned shapes is calculated, and modes of shape variation are computed using PCA.

Each shape sample can then be expressed by means of a set of shape coefficients b_s as follows:

$$x \approx \bar{x} + Q_s \cdot b_s \quad (1)$$

where x is a shape sample, \bar{x} the mean shape, Q_s is the matrix with eigenvectors describing the modes of shape variation in the training set and b_s is the vector containing the coefficients weighting those eigenvectors.

Using the gray value information in the images helps speed up the model search. Therefore, the gray values in the images are warped onto the mean shape, and sampled in an intensity vector g . Similar to the shape model, modes of gray value variation can be calculated by performing a PCA. Each gray value patch can then be expressed by means of a set of gray value coefficients b_g as follows:

$$g \approx \bar{g} + Q_g \cdot b_g \quad (2)$$

where g is a gray value patch, \bar{g} the mean gray value patch, Q_g is the matrix with eigenvectors describing the modes of variation of the gray values and b_g the vector of weight coefficients.

Because shape and gray values are correlated, the eigenvectors describing shape variations and those describing gray value variations can be combined. A PCA can now be performed on the combined vectors and the appearance can be expressed by means of a set of appearance parameters b_a as follows:

$$a \approx \bar{a} + Q_a \cdot b_a \quad (3)$$

where a is an appearance, \bar{a} the mean appearance, Q_a is the matrix with eigenvectors describing the modes of appearance variation and b_a the vector of weight coefficients.

As mentioned above, PCA is used to derive the main directions of shape variation in the dataset. PCA assumes that the input data is drawn from a multidimensional Gaussian distribution, which in practice tends to be an oversimplification. Therefore, we have investigated the use of ICA to find the main directions of shape variation: ICA does not impose normality assumptions in the training data and takes into account higher order moments of the distribution, e.g. skewness and kurtosis. This way, we expect to give a more accurate description of the training data set and achieve better matching results.

3 Independent Component Analysis

3.1 Background

Independent Component Analysis, also known as Blind Source Separation (BSS) [6,9], is a widely used method for separation of mixed signals. This concept can be illustrated as follows. Let X be a shape vector which is assumed to come from a mixture of signals of the form

$$X = A \cdot S \quad (4)$$

where A is a matrix containing mixing parameters and S the source signals/shapes.

The goal of Blind Source Separation is to calculate the original source signals from the mixture. Thus, to estimate the de-mixing matrix U that gives the following:

$$\hat{S} = U \cdot X \quad (5)$$

This method is called blind, because little information is available, i.e. both the mixing matrix A and the matrix containing the sources S are unknown. The de-mixing matrix U is found by optimizing a cost function. Several different cost functions can be used for performing ICA, e.g. kurtosis, neg-entropy, etc. Therefore, different methods exist to calculate U . In previous work, the effect of these methods on the resulting vectors was evaluated [7]. Based on this evaluation, we have chosen to use the JADE method for performing ICA.

JADE, which stands for Joint Approximated Diagonalization of Eigenmatrices [8], is an algorithm based on the joint diagonalization of the cumulant matrices. All cu-

mulants of order 2 and 4 are involved and a joint diagonalization is performed with a Jacobi technique. Because this method does not incorporate a random initialization and is not a gradient-based method, i.e. has no learning rate and other tunable parameters, it is very robust. This method is computationally quite intensive however, because all of the cumulant matrices are diagonalized simultaneously. However, this does not impose any problems, because the training of the model is performed offline. Fig. 1 shows vectors that describe shape variation obtained with ICA and with PCA.




Fig. 1. Mode of shape variation obtained with ICA (left) and mode of shape variation obtained with PCA (right) for 2D short-axis shapes from cardiac MR images. These images show the variation of the ICs and the variation of the eigenvectors when the corresponding weight factors are varied. The ICs show localized shape variations, whereas PCA eigenvectors describe global shape variations.

3.2 Ordering of Independent Components

PCA results in a natural ordering of the vectors according to their variance, because the goal of PCA is to find orthogonal directions along which the projections of the shapes have maximal variance. Eigenvectors with highest variance are placed first and those with low variance are placed last. With ICA, such an ordering is not obtained automatically. However, this ordering is essential for determining which and how many vectors shall be used to give an accurate description of the training dataset: some vectors describe noise effects and are undesired in the shape model. To select the Independent Components that describe “most relevant” independent directions within the training dataset, a sorting method for ICA is required. In [7], a number of sorting methods were evaluated. The method that was found to perform best is explained in the following paragraphs.

When the corresponding weight factor of an IC is varied, the IC shows a variation with respect to the mean with certain amplitude. For noisy modes, these variations are relatively small and not localized. For modes that describe examples as seen in the training data, these variations are localized and have large amplitudes (see Fig. 2).

To quantify the locality of the shape variations and their amplitudes, first the shapes are projected onto each IC. Of these projections a histogram is computed. The width of this histogram, W , is calculated as a measure of variation of the IC. Here, W plays a role similar to the variance in PCA [7].

The shape parameters are varied between limits $-W/2$ and $+W/2$, analogous to PCA where the parameters are varied between -3σ and $+3\sigma$. After varying the parameters within above mentioned limits for each IC, the maximum amplitude of the peaks is determined. Subsequently, to eliminate the effect of noisy peaks, the peaks whose amplitude is larger than 50% of the maximum peak are selected. The total surface O of these peaks is calculated. The ICs are now ordered according to the following criterion:

$$C = W \cdot M \cdot \frac{O}{F} \quad (6)$$

where C is the criterion used for ordering, W is the width of the histogram, M is the maximum amplitude of the largest peak, O is the total surface of the peaks that are larger than 50% of the largest peak and F the ratio between the shape points spanning the peaks and the total number of points in the shape. In this criterion, information about the number of shapes described is included by means of W , the significance of the peak is denoted by M and the ratio (O / F) gives information about the locality of the ICs. Noisy ICs show variations that have small amplitudes and that occur along the entire shape, therefore for these ICs the fraction (O / F) will be small and thus the criterion value will be lower.




Fig. 2. IC describing a mode of shape variation (left) and IC describing noise (right)

3.3 Integration of ICA into AAM-Framework

For the ICA-based AAM, the modes of shape variation in (1) are computed with the JADE method. The modes of gray value variation are computed with PCA. We did not perform ICA on the intensity patches, because we are mainly interested in obtaining a more accurate shape description, and not in a detailed intensity model. The gray value variations only serve as a means to constitute the matching criterion, whereas the shape variations define the allowed deformations. Therefore, the shape part of the model asserts the most critical influence on border localization accuracy.

An appearance model was constructed by combining the eigenvectors for gray value variation with the independent modes of shape variation and performing a PCA on this combination. The appearance model was restricted between -3σ and $+3\sigma$ with the σ resulting from the final appearance PCA.

4 Experimental Validation

4.1 Data

To evaluate the segmentation results obtained with the ICA-based AAM, a set of 150 routinely acquired multislice multiphase short-axis midventricular MR images of the cardiac left and right ventricle in the end-diastolic phase was used. Images were collected from 50 subjects, and the three mid-ventricular slices per subject were used for model training and evaluation. The images had a resolution of 256×256 pixels, a field of view of $400 - 450$ mm and pixel sizes of $1.56 - 1.75$ mm. In each image left- and right-ventricular endo- and epicardial contours were drawn by experts, and were resampled to 105 corresponding landmarks.

4.2 Segmentation results for PCA and ICA

Retaining 98% shape variation resulted in 45 shape eigenvectors in the PCA-based model. To make a fair comparison, ICA-based appearance models were constructed using 45 independent components for shape variation for each model. In addition, appearance models were constructed following the classical approach based on PCA. The computation of the intensity and appearance parts was identical for both model types.

To quantitatively evaluate the ICA-based AAM segmentation, 10 leave-15-out models were trained. This way, 15 samples (5 subjects) were left out and a model was trained on the remaining 135 samples; this was repeated 10 times by leaving out different samples. The performance of the models was tested on the left-out samples.

For each test, the model search was started with the mean model at positions displaced from the known optimal position by ± 10 pixels, following the approach in [10]. The segmentation results were compared on the basis of the average and maximum point-to-point errors between the hand-labeled and the automatically detected contour points (see Table 1). Figure 3 shows representative examples where the ICA-based AAM outperforms PCA-based AAMs. There were no cases where PCA showed better segmentation results than ICA.




Fig. 3. Segmentation results obtained with the PCA-based (top) and ICA-based (bottom) AAM.

Table 1. Comparison of average and maximum point-to-point errors obtained with PCA- and ICA-based Active Appearance Models

| AAM type | Average Point-to-Point Error (pixels) | | | | Maximum Point-to-Point Error (pixels) | | | |
|-----------|---------------------------------------|--------|--------|--------|---------------------------------------|--------|--------|--------|
| | No Offset | | Offset | | No Offset | | Offset | |
| | Median | 90%ile | Median | 90%ile | Median | 90%ile | Median | 90%ile |
| PCA-based | 2.95 | 5.02 | 4.82 | 8.55 | 8.39 | 14.41 | 12.23 | 20.40 |
| ICA-based | 1.86 | 2.84 | 2.78 | 5.34 | 5.01 | 8.18 | 6.90 | 12.57 |

5 Discussion and Conclusions

The results obtained with the ICA-based AAM show a substantially improved accuracy compared to the conventional PCA-based AAM. Because the independent components that describe shape variation allow for localized variations, this model is capable of a much better segmentation. Changing the weight factor of one independent components does not affect the entire shape, which is one of the major advantages of ICA compared to PCA. In general, PCA-based AAMs yield globally good segmentation results. However, as can be seen in Fig. 3, the PCA-based AAM in some cases has great difficulties yielding locally good segmentation results, because the model search in the PCA-based AAM stops when the major part of an image is segmented correctly.

For example, in order to segment the left ventricle correctly, the model has to expand and the right ventricle expands accordingly, and can not be segmented correctly. This way, local segmentation accuracy in the PCA-based model is a tradeoff between right- and left ventricular shape and size. However, Independent Components represent variations in different regions of the contour. Therefore, changing one vector does not affect the others: combining these independent variations proved to yield more accurate segmentation results.

In addition, ICA may give a better description of the input data set, because it does not assume that the data is drawn from a Gaussian distribution and also utilizes higher order moments of the distribution, such as the skewness and kurtosis. This way there is no oversimplification of the data set and a better representation of the shape variations occurring in the training set is obtained.

In conclusion, the presented ICA-based AAM showed excellent border localization accuracy as compared to the classical PCA-based AAM. In the near future, we will further investigate if ICA can give a more compact description of the data, i.e. the possibility to obtain the same accuracy with less independent components to describe shape variations. Also, we will look at the effect of the texture – shape ratio on the segmentation results. Apart from AAMs, we believe ICA is highly promising for use in other statistical shape modeling approaches.

Acknowledgements. M. Üzümcü is supported by grant IBV00303 from IOP, Dutch Ministry of Economic Affairs. A.F. Frangi is supported by a Ramón y Cajal Research Fellowship and grants TIC2002-04495-C02 and TIC2002-10416-E from the Spanish

Ministry of Science and Technology. B.P.F. Lelieveldt is supported by an Innovative Research Incentive 2001 grant from the Netherlands Organization for Scientific Research.

References

1. T.F. Cootes, C. Taylor, A. Hill, J. Halsam, "The Use of Active Shape Models for Locating Structures in Medical Images", *Proc. IPMI 1993*, 33–47, 1993
2. T.F. Cootes, G.J. Edwards, C.J. Taylor, "Active Appearance Models", *Proc. of the ECCV*, vol. 2, 484–498, 1998
3. G. Edwards, C. Taylor, T. Cootes, "Interpreting face images using active appearance Models", *Proceedings of the Third IEEE International Conference on Automatic Face and Gesture Recognition*, 300–305, 1998
4. S. C. Mitchell, B. P. F. Lelieveldt, R. J. van der Geest, J. G. Bosch, J. H. C. Reiber, and M. Sonka, "Multistage Hybrid Active Appearance Model Matching: Segmentation of Left and Right Cardiac Ventricles in Cardiac MR Images," *IEEE TMI*, vol. 20(5), A. pp 415–423, 2001
5. S.C. Mitchell, J.G. Bosch, B.P.F. Lelieveldt, R.J. van der Geest, J.H.C. Reiber, M. Sonka, "3D Active Appearance Models: Segmentation of Cardiac MR and Ultrasound images", *IEEE TMI*, Vol. 21(9), pp 1167–1178, 2002.
6. Hyvärinen and E. Oja, "A Fast Fixed-Point Algorithm for Independent Component Analysis", *Neural Computation*, 9(7), pp. 1483–1492, 1997
7. M. Üzümcü, A.F. Frangi, J.H.C. Reiber, B.P.F. Lelieveldt, "Independent Component Analysis in Statistical Shape Models", *Proc. SPIE 2003*, In Press.
8. J.F. Cardoso, "Higher Order Contrasts for Independent Component Analysis", *Neural Computation*, 11(1), pp. 157–192, 1999
9. Hyvärinen, J. Karhunen and E. Oja, *Independent Component Analysis*, John Wiley & Sons, 2001.
10. T.F. Cootes, P. Kittipanya-ngam, "Comparing Variations on the Active Appearance Model Algorithm", *Proc. BMVC*, Vol.2, pp 837–846, 2002.