General Pose Face Recognition Using Frontal Face Model

Jean-Yves Guillemaut¹, Josef Kittler¹,
Mohammad T. Sadeghi², and William J. Christmas¹

¹ School of Electronics and Physical Sciences, University of Surrey, Guildford, GU2 7XB, U.K.

² Department of Electrical Engineering, Yazd University, Yazd, P.O. BOX 89195-741, Iran

Abstract. We present a face recognition system able to identify people from a single non-frontal image in an arbitrary pose. The key component of the system is a novel pose correction technique based on Active Appearance Models (AAMs), which is used to remap probe images into a frontal pose similar to that of gallery images. The method generalises previous pose correction algorithms based on AAMs to multiple axis head rotations. We show that such model can be combined with image warping techniques to increase the textural content of the images synthesised. We also show that bilateral symmetry of faces can be exploited to improve recognition. Experiments on a database of 570 non-frontal test images, which includes 148 different identities, show that the method produces a significant increase in the success rate (up to 77.4%) compared to conventional recognition techniques which do not consider pose correction.

1 Introduction

Face recognition has been a topic of active research in computer vision and pattern recognition for several decades. Applications encompass many aspects of everyday life such as video surveillance, human-machine interface or multimedia applications. Some of the reasons why face recognition has been attracting so much attention is that, unlike other biometrics such as fingerprints or eye iris scan, it does not require cooperation of the subject, it is unobtrusive and it can be done with a relatively cheap equipment. Despite the high recognition rates achieved by current recognition systems in the case of frontal images, performance has been observed to drop significantly if such ideal conditions are not satisfied. In fact, a previous evaluation of face recognition algorithms [1] has identified the face recognition problem from non-frontal images as a major research issue. In this paper, we concentrate on this problem and propose a novel solution.

1.1 Previous Work

Early solutions to the general pose face recognition problem were multi-view generalisations of standard frontal face recognition techniques. In [2], Beymer extends template-based techniques to non-frontal poses by building galleries of views for some pose configurations which sample the viewing sphere. In [3], Pentland et al. apply the eigenspace technique to arbitrary pose images by building a separate eigenspaces for

each pose configuration. One major limitation of such methods is that a large number of images is required to sample the viewing sphere for each subject.

More recent work focused on eliminating the effects of pose variation by remapping gallery and probe images into similar pose configurations in which case standard recognition techniques are known to perform well. In [4], Vetter and Poggio show that such image transformations can be learnt from a set of prototypical images of objects of the same class that form what they call Linear Object Classes. They synthesised realistic frontal images of faces from non-frontal views, however the decomposition as a sum of basis functions results in a loss of textural information. In [5] Vetter addresses this problem by supplementing the previous method with a generic 3D model which remaps the texture of the original image onto the corrected view. One limitation is that a database of prototype images is needed for each pose that must be corrected or synthesised, which requires the acquisition of a large number of images.

A different line of research concerns the use of parametric surfaces in the recognition feature space. The principle has been formulated in a general object recognition context in [6]. In this work Murase and Nayar consider the set of images of an object undergoing a rotation motion in 3D space and subject to changes of illumination. They observed that the projection of such images into the eigenspace forms a characteristic hypersurface for each object. The recognition problem is then reformulated in terms of finding the hypersurface which lies closest to the projection of the probe image in the eigenspace for a given metric. The principle has been applied in face recognition in the case of single [7] or multiple [8] images. A major limitation of such methods is that the construction of the eigensurface requires a large number of images for each subject.

Another important class of methods consists of model-based methods. The general idea is that a face in an image can be represented by the parameters of a model which can be used for recognition. In [9], Wiskott et al. represent faces by labelled graphs, where each node is labelled with a set of complex Gabor wavelet coefficients, called a jet. In [10], an Active Appearance Model (AAM) [11] is used for face localisation and recognition. The authors used Linear Discriminant Analysis to separate the parameters encoding the identity from the parameters encoding other sources of variation (pose, illumination, expression). In [12], the authors show that the appearance parameters can be used to estimate the pose of a face and synthesise novel views in different poses. They apply the method successfully to tracking images in [12] and face verification [13]. 3D morphable models have also been used to localise faces and identify subjects based on the fitted model parameters [14] or a corrected frontal view [15]. 3D morphable models handle better occlusions than AAM, however they require better initialisation and their convergence may be more difficult.

Finally, in [16] a view synthesis technique based on shape-from-shading is used to correct images with arbitrary poses and illumination into a frontal view under frontal lighting. Unlike other methods, this approach does not require a large number of example images, however light source and pose estimation were done manually.

1.2 Our Approach

Our approach is based on using an AAM to localise the face and synthesise a frontal view which can be then fed into a conventional face recognition system. We require

only a gallery of frontal views of each subjects (e.g. mugshots) to train the recognition system, and we use only a single image of a subject in an arbitrary pose (usually non-frontal) for identification. This is a significant advantage compared to techniques requiring multiple example views for training [2,3,7,8]. Another strong point of our system is that it has the potential to localise automatically facial features in the image. This contrasts with a number of approaches which rely more heavily on a good initialisation [14,15,16].

Our approach is different from previous AAM-based face recognition systems [10,12,13] in the sense it does not use the model parameters for recognition. Instead it uses a corrected frontal appearance whose shape is predicted by a statistical model, modelling pose variation, and whose texture is either directly synthesised from the appearance model or obtained by image warping techniques. The latter approach presents the advantage of preserving the textural information (moles, freckles, etc) contained in the original image; such information would be lost in a traditional model parameter representation which models only the principal components of the appearance (low-pass filter equivalent). Another specificity of our pose correction model is that it can accommodate more general head rotations than the original model [12] which was formulated for single axis rotation only.

Our main contributions are the following. Firstly we formulate a novel pose correction method based on AAMs which generalises previous methods [12] as described in the previous paragraph. Secondly we show that AAMs can be used to improve face recognition performance by synthesis of corrected views of the probe images. Finally, we show that the bilateral symmetry of the face can be exploited to attenuate the effect of occlusions and increase the recognition performance.

The paper is structured as follows. We start by giving an overview of the system. We then concentrate on the novel pose estimation and correction algorithm proposed. Experimental results are given on a database of non-frontal images.

2 Methodology

The system is illustrated in Fig. 1. We give a brief description of each module.

Fig. 1. Illustration of the main modules constituting the system

Face localisation. An AAM [11] is used for localising faces and their characteristic features. Our implementation uses 128 feature points for shape description. For efficiency, a multi-resolution approach with three different resolutions is adopted. The AAM is initialised with the coordinates of the eye centres, which could be obtained for example from an eye detector. In order to improve the convergence properties of algorithm in the case of non-frontal images, we use five different initialisations corresponding to mean appearance for different poses and select the result with lowest residual.

Pose estimation and pose correction. The aim of these modules are firstly to estimate the pose of the face in the probe image and then to synthesise a novel view of the subject in the same pose as the gallery images, i.e. frontal in this case. This is the core of our method. It will be described in detail in the next section.

Geometric and photometric normalisation. Geometric normalisation is done by applying an affine transformation composed of a translation, rotation and scaling in order to align the eye centres with some pre-defined positions; the position of the eye centres in the original image is obtained automatically from the fitted appearance. Photometric normalisation is done by histogram equalisation [17].

Identification. The statistical features used for recognition are obtained by Linear Discriminant Analysis (LDA) [18]. Identification is done by comparing the projection of the probe and gallery images in the LDA subspace and selecting the gallery image which maximises the normalised correlation [19]. Our implementation uses the bilateral symmetry of faces to attenuate the effect of occlusions (see details in result section).

3 Pose Estimation and Correction

Our method is inspired from the work of Cootes et al. described in [12]. In this paper, the authors formulated a pose correction method which handles rotation around a unique axis. Although it was claimed that generalisation to more general rotations was straightforward, no explicit formulation was given. In [13], it was suggested that rotation around two axes could be handled by using sequentially two independent pose correction models trained for pan and tilt motion respectively. Although this may work in practice for small rotations, this is not suitable for correcting rotations which exhibit simultaneously large pan and tilt components because such poses have not been learnt by either pose correction model. We formulate a pose correction method which handles correctly simultaneous pan and tilt head rotations. In addition, we show that image warping techniques can be used to improve the textural content of the corrected images.

3.1 Modelling Pose Variation

Out of plane head rotation is parametrised by two angles: the pan angle θ and the tilt angle ϕ , accounting respectively for rotation around the vertical axis and the horizontal axis attached to the face. This is sufficient to parametrise arbitrary head pose, because in-plane rotation, translation and image scaling are already modelled by the appearance model parameters. In an appropriately defined world reference frame, a feature point

attached to the head and with coordinates $(X_0, Y_0, Z_0)^\top$ transforms into the point with coordinates $(X, Y, Z)^\top$ after a rotation parametrised by (θ, ϕ) , such that:

$$X = X_0 c_\theta c_\phi - Y_0 s_\theta c_\phi + Z_0 s_\phi, \quad Y = X_0 s_\theta + Y_0 c_\theta, \quad \text{and} \quad Z = -X_0 c_\theta s_\phi + Y_0 s_\theta s_\phi + Z_0 c_\phi, \quad (1)$$

where we use the notations $c_\alpha = \cos \alpha$ and $s_\alpha = \sin \alpha$. Assuming an affine projection model, the 3D point $(X, Y, Z)^\top$ projects into the image point $(x, y)^\top$ such that:

$$x = x_0 + x_1 c_\theta + x_2 s_\theta + x_3 c_\phi + x_4 s_\phi + x_5 c_\theta c_\phi + x_6 c_\theta s_\phi + x_7 s_\theta c_\phi + x_8 s_\theta s_\phi, \quad (2)$$

where $x_0, \dots, x_8$ are some constant (a similar equation is obtained for y). The shape model being linear, the shape parameters c follow a similar linear model:

$$c = c_0 + c_1 c_\theta + c_2 s_\theta + c_3 c_\phi + c_4 s_\phi + c_5 c_\theta c_\phi + c_6 c_\theta s_\phi + c_7 s_\theta c_\phi + c_8 s_\theta s_\phi, \quad (3)$$

where $c_0, \dots, c_8$ are constant vectors which can be learnt from a database of annotated images. Experiments we carried out suggest that this equation can be extended to the appearance parameters. This is consistent with what was observed by Cootes et al. in the case of a single rotation in [12]. Note that if one of the angles is set to a fixed value, (3) simplifies to the equation originally formulated in [12].

3.2 Pose Estimation

We define the matrix $R_c = [c_1 | c_2 | c_3 | c_4 | c_5 | c_6 | c_7 | c_8]$. Given a vector c of shape or appearance parameters, we compute the vector $[a_1, \dots, a_8]^\top = R_c^+(c - c_0)$, where R_c^+ is the pseudo-inverse of R_c . A closed-form solution for the pan and tilt angles is then given by: $\theta = \tan^{-1}(\frac{a_2}{a_1})$ and $\phi = \tan^{-1}(\frac{a_4}{a_3})$. Such a solution is not optimum because it involves only the values a_1 to a_4 . A more accurate solution is obtained by finding the values of θ and ϕ which minimise the following cost function:

$$d_c(\theta, \phi) = \|c - (c_0 + c_1 c_\theta + c_2 s_\theta + c_3 c_\phi + c_4 s_\phi + c_5 c_\theta c_\phi + c_6 c_\theta s_\phi + c_7 s_\theta c_\phi + c_8 s_\theta s_\phi)\|. \quad (4)$$

This is a simple two-dimensional non-linear minimisation problem which can be solved e.g. with a steepest descent algorithm initialised with the closed-form solution.

3.3 Synthesising Corrected Views

We assume that the pose in the original image has been estimated as (θ, ϕ) and would like to synthesise a novel view of the same subject in the pose (θ', ϕ') . As in [12], we compute the residual vector r not explained by the pose model in the original image:

$$r = c - (c_0 + c_1 c_\theta + c_2 s_\theta + c_3 c_\phi + c_4 s_\phi + c_5 c_\theta c_\phi + c_6 c_\theta s_\phi + c_7 s_\theta c_\phi + c_8 s_\theta s_\phi). \quad (5)$$

The shape or appearance parameters c' of the rotated view in the new pose (θ', ϕ') are then obtained by re-injecting the residual vector r into the new pose equation:

$$c' = c_0 + c_1 c'_{\theta} + c_2 s'_{\theta} + c_3 c'_{\phi} + c_4 s'_{\phi} + c_5 c'_{\theta} c'_{\phi} + c_6 c'_{\theta} s'_{\phi} + c_7 s'_{\theta} c'_{\phi} + c_8 s'_{\theta} s'_{\phi} + r. \quad (6)$$

If (6) is applied to all appearance parameters, the appearance model can be then used to synthesise a full corrected view of the person (see second row in Fig. 2). We will refer to this method as the *basic pose correction* method.

Fig. 2. Example of non-frontal images (top row) and corrected frontal images (middle rows). For comparison, the bottom row shows example of real frontal images of the same subjects.

3.4 Improving the Textural Content of the Corrected Views

The novel view synthesis method described in the previous section solves elegantly the pose correction problem by predicting the appearance parameters of the novel view and then synthesising a full appearance. There are however two limitations to this approach. Firstly, details such as e.g. moles or freckles are lost in the corrected view, because the appearance parameter representation preserves only the principal components of the image variations. Another limitation is that the basic pose correction method is able to predict the appearance only within the convex hull of the set of feature points, which explains why a black border is present around the face. In practice, this may pose problem during recognition if such border is not present in the gallery images.

We present two methods based on image warping which do not suffer from such limitations. The key idea is to apply (6) only to the shape parameters. This yields an estimate of the position of the feature points in a frontal view. Then the texture of the corrected image is obtained by warping the original image. Two warping techniques have been considered: i) piece-wise affine warping and ii) thin-plate spline warping. Results for all methods are illustrated in Fig. 2 for a few randomly selected subjects. In the first approach, meshes of the original and corrected faces are generated, with vertices placed at the localised or predicted feature points. Triangular meshes are generated automatically by Delaunay triangulation [20]. Then each triangle is warped affinely to its new position (see Fig. 3). The second technique is based on thin-plate splines [21]. It has the advantage of resulting in smoother deformations than the previous method (no

artefact at the boundary between triangles), however the behaviour of the method is not always clear in-between feature points, especially in the case of large pose variations.

Experiments carried out on a database of 396 images of unknown subjects (not used for training the pose correction model) with variations of $\pm 22.5^\circ$ for the pan angle and $\pm 30^\circ$ for the tilt angle showed that our pose estimation model is accurate to about 5° for the pan angle and 7° for the tilt angle. Errors in pose estimation translate into errors in the computation of r in (5), which ultimately result in errors in the corrected frontal view. Typically, errors in estimation of pan and tilt angle result in a compression or expansion of the face in the horizontal or vertical direction respectively. Errors in scale in the horizontal direction are usually less problematic because the distance between the two eyes is normalised for recognition. Unfortunately there exists no such compensation for scaling errors in the vertical direction.

Fig. 3. Illustration of the piece-wise affine warping method. The original image (arbitrary pose) is shown on the left, while the corrected image (frontal pose) is shown on the right. Both images are overlaid with the triangular mesh used for warping.

4 Experimental Results

Experiments were carried out on the XM2VTS database [22]. The database contains eight frontal images and four non-frontal images (corresponding to poses with head turned left, right, up or down) of 295 different subjects. Among all these images, 225 frontal images and 1177 non-frontal images have been manually annotated with facial landmarks. We also use a database of 567 similarly annotated images of 43 subjects for which additional pose information is available. Ground truth pose (pan and tilt angles) was obtained by placing the subjects on a turntable during image acquisition.

For a fair evaluation, the images have been split into two subsets. The first subset contains the images of the first 147 subjects from the XM2VTS database plus the turntable images, and has been used to train the AAM and the pose correction model (when pose ground truth was available). The images of the remaining 148 subjects from the XM2VTS database (570 images in total) are used for recognition experiments; the frontal images were used for training the recognition system (gallery images), while the non-frontal images were used for testing (probe images). None of the subjects used for training the AAM or the pose correction model were used during the recognition experiments.

Two different test sets were considered. *Test set 1* (295 images) contains only the probe images for which the subjects have their eyes open and do not wear glasses. Eyes closed or glasses (which can generate specularities) complicate significantly the problem because the eyes, which contain important information for identification and face localisation, may not be visible. *Test set 2* contains all probe images (570 images). Both test sets are very challenging because of the large pose variations observed (see top row of Fig. 2 for some example of probe images).

Experiments were carried out in two modes: *manual* and *semi-automatic*. In the manual mode, the system is initialised with the manually marked-up feature points; this eliminates potential errors due to AAM fitting and allows us to measure the performance of the system independently from the localisation algorithm. In the semi-automatic mode, faces are localised by the AAM algorithm initialised with the coordinates of the eye centres obtained from manual annotation. In future implementation, the method will be made fully automatic by using an eye detector to initialise the AAM search.

Four different methods are compared. The method with *no pose correction* applies only geometric and photometric normalisation to the test images before projection into the LDA subspace. For geometric normalisation, the images are cropped to a window of dimension 55×51 pixels, where the left and right eyes occupy the points with coordinates (19, 38) and (19, 12). This is a conventional recognition method which is known to perform well in the case of frontal images. The other methods apply additional pose correction based techniques described earlier: *basic pose correction* (see Sect. 3.3), shape correction combined with either *piece-wise affine warping* or *thin-plate spline warping* (see Sect. 3.4).

Given the large changes of pose observed in the images, parts of the face can become largely occluded, which can produce significant artefacts in the corrected images. In order to attenuate such effects, at least in the case of rotations around the vertical axis, the bilateral face symmetry has been used to eliminate the occluded half of the face when needed. In this approach, three different LDA subspaces are build for full image, left half-image and right half-image respectively. Then the pose estimate for the probe images is used to select automatically the most appropriate LDA subspace to use for identification. At the moment, the pose classification is done by thresholding of the pan angle (thresholds of -15° and $+15^\circ$ have been used). All recognition methods are tested with and without this bilateral face symmetry based occlusion removal algorithm; we refer to these methods as *partial face* and *full face* methods respectively.

The success rates (percentage of subjects identified as top matches) obtained for each configuration are shown in Table 1. The best performing method is the one which uses shape correction combined with a piece-wise affine warping, followed very closely by shape correction combined with the thin-plate spline warping. Compared to a conventional face recognition system which does not consider pose correction, the best pose correction method improves the success rate by between 33.7% and 77.4% depending on the difficulty of the test set and the degree of initialisation. The best success rate measured is 69.2%. This is a high recognition score given the number of classes (148 subjects) and the fact that all images are non-frontal (pure chance would be only 0.67%). The basic pose correction method is the least accurate. This suggests that it is important to preserve the textural information contained in original images. The loss of information in the image synthesised from the predicted frontal appearance parameters is accentuated by errors in locating the face in the case of the semi-automatic algorithm. It can be observed that the use of bilateral face symmetry for reducing the effect of occlusions allows to increase the performance by a few percents in the case of the semi-automatic algorithm; it is not as critical in the case of manually localised faces.

Table 1. Success rate for different general pose face recognition methods

		no pose correction		basic pose correction		piece-wise affine warping		thin-plate spline warping	
		full face	part. face	full face	part. face	full face	part. face	full face	part. face
Test set 1	manual	39.0	38.0	33.2	32.2	69.2	69.2	66.1	66.8
	semi-auto	39.0	38.0	17.3	16.6	60.0	63.7	56.9	59.7
Test set 2	manual	40.4	40.0	33.5	33.3	62.6	62.6	59.3	59.3
	semi-auto	40.4	40.0	13.7	13.7	52.5	54.0	50.2	51.9

5 Conclusions and Future Work

We have presented a novel method for face identification which is able to cope with pose variations and requires only a single view of a person in an arbitrary pose for identification. The method relies on the use of a statistical model to estimate and synthesise frontal views of the subjects. When combined with image warping techniques, the method is able to preserve the textural content of the original non-frontal image. The corrected image can be fed directly into a conventional face recognition system. It has been shown that such a correction algorithm is able to improve the performance by up to 77.4% compared to a conventional approach which does not consider correction. We also showed how bilateral face symmetry can be used to alleviate the effects of occlusions by using the pose estimate to classify images into three categories for which separate LDA subspaces have been built.

We have compared several methods for correcting the pose and applied them successfully to the problem of face recognition. We are currently working on comparing these methods with other approaches which carry out the recognition directly in the space of model parameters after having decoupled the parameters encoding the identity from the ones encoding pose, expression and illumination [10]. Although the comparison is still in its early stages, we can already anticipate that such method will probably not be able to achieve as high success rates as the ones given here because of the loss of texture information induced by the model parameter representation.

We think that there is a scope for improving further the technique presented in this paper. One possible avenue for future work is to investigate how pose estimation (and thereby pose correction) can be improved by treating the problem jointly with the face recognition problem; in this approach an optimum pose estimate is found by minimising the metric used for matching in the LDA subspace. Other possible avenues include the use of non-linear techniques such as kernel PCA to improve the performance of our AAM in the case of pose variation, a better handling of occlusions (at the moment we classify faces in only three classes according to pan angle) or the extension of the method to image sequences.

Acknowledgements

This work was initiated under the EU Project VAMPIRE and is now supported by the EU Network of Excellence BIOSECURE, with contributions from EU Project MUSCLE and EPSRC Project 2D+3D=ID.

References

1. Phillips, P., Grother, P., Micheals, R., Blackburn, D., Tabassi, E., Bone, J.: Face recognition vendor test 2002: Evaluation report (2003)
2. Beymer, D.J.: Face recognition under varying pose. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition. (1994) 756–761
3. Pentland, A., Moghaddam, B., Starner, T.: View-based and modular eigenspaces for face recognition. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition. (1994) 84–91
4. Vetter, T., Poggio, T.: Linear object classes and image synthesis from a single example image. IEEE Transactions on Pattern Analysis and Machine Intelligence **19**(7) (1997) 733–742
5. Vetter, T.: Synthesis of novel views from a single face image. International Journal of Computer Vision **28**(2) (1998) 103–116
6. Murase, H., Nayar, S.K.: Visual learning and recognition of 3-d objects from appearance. International Journal of Computer Vision **14**(1) (1995) 5–24
7. Graham, D.B., Allinson, N.M.: Face recognition from unfamiliar views: Subspace methods and pose dependency. In: Proc. IEEE International Conference on Automatic Face and Gesture Recognition. (1998) 348–353
8. Li, Y., Gong, S., Liddell, H.: Constructing facial identity surfaces for recognition. International Journal of Computer Vision **53**(1) (2003) 71–92
9. Wiskott, L., Fellous, J.M., Kruger, N., von der Malsburg, C.: Face recognition by elastic bunch graph matching. IEEE Transactions on Pattern Analysis and Machine Intelligence **19**(7) (1997) 775–779
10. Edwards, G.J., Cootes, T.F., Taylor, C.J.: Face recognition using active appearance models. In: Proc. European Conference on Computer Vision. Volume II. (1998) 581–595
11. Cootes, T.F., Edwards, G.J., Taylor, C.J.: Active appearance models. In: Proceedings of European Conference on Computer Vision. Volume II. (1998) 484–498
12. Cootes, T.F., Wheeler, G.V., Walker, K.N., Taylor, C.J.: View-based active appearance models. Image and Vision Computing **20**(9-10) (2002) 657–664
13. H. Kang, T.F.C., Taylor, C.J.: A comparison of face verification algorithms using appearance models. In: Proc. British Machine Vision Conference. Volume 2. (2002) 477–486
14. Blanz, V., Vetter, T.: Face recognition based on fitting a 3d morphable model. IEEE Transactions on Pattern Analysis and Machine Intelligence **25**(9) (2003) 1063–1074
15. Blanz, V., Grother, P., Phillips, P.J., Vetter, T.: Face recognition based on frontal views generated from non-frontal images. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition. Volume 2. (2005) 454–461
16. Zhao, W., Chellappa, R.: SFS based view synthesis for robust face recognition. In: Proc. IEEE International Conference on Automatic Face and Gesture Recognition. (2000) 285–292
17. Gonzalez, R.C., Woods, R.E.: Digital Image Processing. Second edn. Prentice Hall (2002)
18. Devijver, P.A., Kittler, J.: Pattern Recognition: A Statistical Approach. Prentice-Hall (1982)
19. Li, Y.: Linear Discriminant Analysis and its application to Face Identification. PhD thesis, University of Surrey (2000)
20. de Berg, M., Schwarzkopf, O., van Kreveld, M., Overmars, M.: Computational Geometry: Algorithms and Applications. second edn. Springer-Verlag (2000)
21. Bookstein, F.L.: Principal warps: Thin-plate splines and the decomposition of deformations. IEEE Transactions on Pattern Analysis and Machine Intelligence **11**(6) (1989) 567–585
22. Messer, K., Matas, J., Kittler, J., Luettin, J., Maitre, G.: XM2VTSDB: The extended M2VTS database. In: Proceedings of International Conference on Audio- and VideoBased Biometric Person Authentication (AVBPA). (1999) 72–77