Robust Face Alignment for Illumination and Pose Invariant Face Recognition

Fatih Kahraman[1], Binnur Kurt[2], Muhittin Gökmen[2]
Istanbul Technical University, [1]Informatics Institute, [2]Computer Engineering Department
34469 Ayazağa, İstanbul, Turkey
kahraman@be.itu.edu.tr, {gokmen, bkurt}@itu.edu.tr

Abstract

In building a face recognition system for real-life scenarios, one usually faces the problem that is the selection of a feature-space and preprocessing methods such as alignment under varying illumination conditions and poses. In this study, we developed a robust face alignment approach based on Active Appearance Model (AAM) by inserting an illumination normalization module into the standard AAM searching procedure and inserting different poses of the same identity into the training set. The modified AAM search can now handle both illumination and pose variations in the same epoch, hence it provides better convergence in both point-to-point and point-to-curve senses. We also investigate how face recognition performance is affected by the selection of feature space as well as the proposed alignment method. The experimental results show that the combined pose alignment and illumination normalization methods increase the recognition rates considerably for all feature-spaces.

1. Introduction

Face recognition systems have matured from the systems working only in highly controlled indoor environments to the systems allowing identification of individuals in indoor or outdoor environments under severe conditions. But some problems still remain, constraining their success to a limited degree. Largely illumination and pose variations are responsible for dramatic variations on the appearance of the same individual. Hence any improvement in face appearance will enhance the recognition performance. These variations lead to complex effects imposed on the acquired face image that pertains little to the actual identity. Face recognition systems are usually required to handle highly varying illumination and pose conditions. As face recognition techniques advance, more researchers have focused on solving issues arising from illumination and pose in one shot.

Face alignment is a very important step to extract good facial features to obtain high performance in face recognition, expression analysis and face animation applications. Several face alignment methods were proposed: Active Shape Models (ASM) [1] and Active Appearance Models (AAM) [2] [3], proposed by Cootes et al are two successful models for object localization. ASM uses local appearance models to find the candidate shape and global model to constrain the searched shape. AAM combines the constraints on both shape and texture variations in its characterization of facial appearance. In searching for a solution, it assumes linear relationships between appearance variation and texture variation and between texture variation and position variation. In this study, we have used AAM to solve the pose-invariant face alignment problem.

Image variation due to lighting changes is larger than that due to different personal identities. Because lighting direction changes alter the relative gray scale distribution of face image. Consequently, illumination normalization is required to reach acceptable recognition rates. Varying illumination is a difficult problem and has received much attention in recent years. Two studies among them are very important: symmetric shape from shading [4] and illumination cones [5] where face image variations due to light direction changes are theoretically explained. In the later algorithm, both self shadow and cast-shadow were considered and its experimental results outperformed most of the existing methods. The major drawbacks of the illumination cone model are the computational cost and the strict requirement of seven input images per person. Basri et al [6] represent lighting using a spherical harmonic basis wherein a low-dimensional linear subspace is shown to be quite effective for recognition. The harmonic images can easily be computed analytically given surface normals and the albedos. Shashua [7] employ a very simple and practical image ratio method to map the face images into different lighting conditions. There are several recent image-based studies on illumination invariant face recognition. Image-based methods are known to be robust to illumination variations [8]. Main drawback of the image-based methods is that they always assume the face image is already aligned. Usually it is not an easy assumption to satisfy especially when the input image is poorly illuminated. AAM is known to be very sensitive to illumination, particularly if the lighting conditions during testing are significantly different from the lighting conditions during training. Several variations of AAM appear in the literature to

improve the original algorithm, namely view-based AAM [9], Direct Appearance Models [10]. Despite the success of these methods, problems still remain to be solved. Moreover, under the presence of partial occlusion, the PCA-based texture model of AAM causes the reconstruction error to be globally spread over the image, thus degrading alignment. In this paper, we propose an approach based on histogram-fitting to overcome the problem explained above. A detailed explanation of the proposed approach is given in Section.2.

Yet another issue related to face recognition is to recognize different poses of the same person. Pose-invariant face recognition requires pose alignment where images are either captured by multiple cameras or by a single camera at different time instances. There are several works related to pose normalization. Blanz and Vettel [11] use a statistical 3D morphable model to tackle with pose and illumination variations. Since their method requires textured 3D scans of heads, it is computationally expensive. Cootes et al constructed three AAMs which are called as View-based AAMs [9]. These models are linear model of frontal, profile and half profile views of faces. They also show how to estimate the pose from the model parameters. The approach in this study differs from their method in the way that we construct only one AAM rather than three models. The idea here is to reduce three searching procedures to only one fitting procedure by using one statistically powerful model to generalize pose variations. In order to do that we trained one linear model by using a training dataset consisting of 8 different poses of 3 individuals captured under similar illumination conditions. In Section 3, we will study the AAM capable of producing different poses of unseen person and show how we project a non-frontal face to a frontal face.

In this paper, we focus on the problems induced by varying illumination and poses. Our primary aim is to eliminate the negative effect of illumination and pose on the face recognition system performance through illumination and pose-invariant face alignment based on Active Appearance Model. The rest of the paper is structured as follows: Section 2 introduces Active Appearance Model (AAM) and illumination normalization inserted into the searching procedure of AAM. Section 3 is for the proposed pose invariant combined active appearance model. The experimental results and the conclusion are presented in Section 4 and 5, respectively.

2. Active Appearance Model

Active Appearance Models (AAM) are generative models capable of synthesizing image of a given object class. By estimating a compact and specific basis from a training set, model parameters can be adjusted to fit unseen images and hence perform image interpretation. The modeled object properties are usually shape and pixel intensities. Training objects are defined by marking up each example image with point of correspondence. AAMs can be rapidly fitted to unseen images, given a reasonable initialization.

AAM works according to the following principle: A face image is marked with *n* landmark points. The content of the marked face is analyzed based on a Principal Component Analysis (PCA) of both face texture and face shape. Face shape is defined by a triangular mesh and the vertex locations of the mesh. Mathematically the shape model is represented as follows: $x = [x_1, x_2, \ldots, x_n, y_1, y_2, \ldots, y_n]$. Face texture is the intensities on these landmarks (color pixel values normalized to shape) and is represented with the formula (g). Face shape and texture are reduced to a more compact form through PCA such that $x = \bar{x} + \Phi_s b_s$ and $g = \bar{g} + \Phi_g b_g$. In this form, Φ_s contains the t eigenvectors corresponding to the largest eigenvalues and b_s is a t-dimensional vector. By varying the parameters in b_s, the shape can be varied. In the linear model of texture, Φ_g is a set of orthogonal modes of variation and b_g is a set of grey-level parameters. To remove the correlation between shape and texture model parameters, a third PCA is applied on the combined model parameters such that $x = \bar{x} + \Phi_s W_s^{-1} Q_s c$ and $g = \bar{g} + \Phi_g Q_g c$ where $b = \begin{bmatrix} W_s b_s & b_g \end{bmatrix}^T$ and $b = \begin{bmatrix} Q_s & Q_g \end{bmatrix}^T c$.

In this form, W_s is a diagonal matrix of weights for each shape parameter, allowing for the difference in units between the shape and the grey models; c is a vector of appearance parameters controlling both the shape and the grey-levels of the model. Q_s and Q_g are the eigenvectors of the shape and texture models respectively.

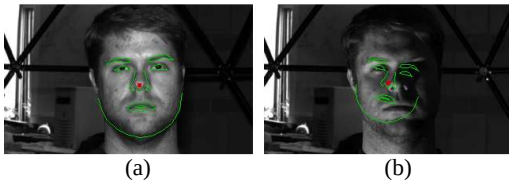

(a) (b)

Figure.1 Face alignment using standard AAM under good and extreme illumination. (a) Normal illumination, (b) Extreme illumination

We propose an illumination normalization method in order to increase the accuracy of AAM applied to images captured under different illumination conditions by inserting an illumination normalization module into the standard AAM searching procedure. The problem is demonstrated in Fig.1. In Fig.1 (a) a correct AAM search result is shown where the input image contains a frontal face illuminated frontally.

2.1. Illumination Normalization

We discuss here two light normalization methods and we analyze their behavior when used in AAM searching. The first proposed method is ratio-image [12] face illumination normalization method. Ratio-image is defined as the quotient between a face image whose lighting condition is to be normalized and a reference face image. These two images are blurred using a Gaussian filter, and the reference image is then updated by an iterative strategy in order to further improve the quality of the restored face. Using this illumination restoration method, a face image with arbitrary illumination can be restored to a face having frontal illumination.

The second normalization method discussed in this study is based on image histogram techniques. The global histogram equalization methods used in image processing for normalization only transfers the holistic image from one gray scale distribution to another. This processing ignores the face-specific information and cannot normalize these gray level distribution variations. To deal with this problem, researchers have made many improvements in recent years. The problem is that well-lit faces do not have a uniform histogram distribution and this process gives rise to an unnatural illumination to the face. As suggested in [13], it is possible to normalize a poorly illuminated image via histogram fitting to a similar, well illuminated image. In this study we used a special type of histogram fitting algorithm for face illumination normalization.

We make our analysis on one particular case where one side of the face is dark and the other side is bright. The main idea here is to fit the histogram of the input face image to the histogram of the mean face. The face is first broken into two parts (left/right) and then the histogram of each window is independently fitted to the histogram of mean face. For these two histograms, namely the histogram of the left window denoted as $H_l(i)$ and the histogram of the right window denoted as $H_r(i)$, two mapping functions are computed: $f_{H_l \to G}$ and $f_{H_r \to G}$ corresponding to the left and right windows. Here $G(i)$ is the histogram of the reference image also called mean face in AAM. An artifact introduced by this mapping is the sudden discontinuity in illumination as we switch from the left side of the face to the right side. The problem is solved by averaging the effects of the two mapping functions with a linear weighting that slowly favors one for the other as we move from the left side to the right side of the face. This is implemented with the mapping function $f_{H_{total} \to G}$ defined as bellow:

$$f_{H_{total} \to G}(i) = leftness \times f_{H_l \to G}(i) + (1 - leftness) \times f_{H_r \to G}(i).$$

Lighting normalization result is shown in Fig.2 obtained by using the histogram fitting method explained above together with the histogram plots. Fig.2 (a) shows the reference image and its histogram. The normalization algorithm is applied to the input image given in Fig.2 (b). The normalization result is shown in Fig.2 (c) where the histogram of the restored image is very close to the histogram of the reference image as expected.

As it can be seen from Fig.3 and Fig.4 the normalization method can produce more suitable images to be used in AAM search mechanism. The classical AAM search fails for all images given in the first row of Fig.3. We will show in the next section that AAM search procedure can now converge to the correct shape for the restored image both in point-to-point error and point-to-curve sense. Fig.4 presents several results obtained for Set 4 (left) and Set 3 (right) faces of different individuals having extremely dark and bright regions. A significant amount of improvement in the quality can be easily verified from the experimental results. The dark parts now become somehow noisy whereas there are still some very bright areas.

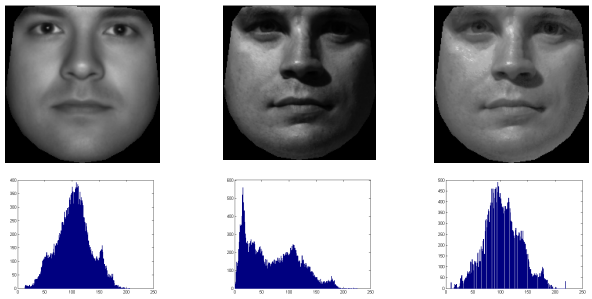

(a) (b) (c)

Figure.2 Light normalization using histogram fitting**:** (a) Mean face and its histogram, (b) Test face and its histogram, (c) Normalized face and its histogram.

Figure.3 Light normalization results: On the top the input images are given, and on the bottom the normalized images are shown.

2.2. Experimental Results

AAM combines the shape and texture model in one single model. The alignment algorithm (also called AAM searching) optimizes the model in the context of a test image of a face. The optimization criterion is the error between a synthesized face texture and the corresponding texture of the test image.

Figure.4 Light normalization results for extreme cases: On the top the input images are given, and on the bottom the normalized images are shown.

Due to the illumination problems the error can be high and the classic searching algorithm fails. In the proposed approach, we normalize the corresponding texture in the test image just before we compute the error. We tested the proposed method on the Yale-B [14] face dataset. The total number of images under different lighting conditions for each individual is 64. The database is portioned into four sets identified as Set 1-4. Set 1 contains face images whose light direction is less than ±20 degrees. Set 2 contains face images whose light directions are between ±20 and ±50 degrees. Set 3 contains face images whose light directions are between ±50 and ±70 degrees. Set 4 contains face images whose light directions are greater than ±70 degrees. All details about the Yale B dataset are given in [14]. We manually labeled 4920 images. To establish the models, 73 landmarks were placed on each face image; 14 points for mouth, 12 points for nose, 9 points for left eye, 9 points for right eye, 8 points for left eyebrow, 8 points for right eyebrow and 11 points for chin. The warped images have approximately 32533 pixels inside the facial mask. We constructed a shape space to represent 95% of observed variation. Then we warped all images into the mean shape using triangulation. Using normalized textures, we constructed a 21-dimensional texture space to represent 95% of the observed variation in textures and for shapes we constructed a 12-dimensional shape space to represent 95% of the observed variation in shapes. Finally, we constructed a 15-dimensional appearance space to represent 95% of the total variation observed in the combined (shape and texture) coefficients.

Using a ground truth given by a finite set of landmarks for each example, performance can be easily calculated. In a leave-one-out setting, a distance measure, $D(x_{gt},x)$, is computed that gives a scalar interpretation of the fit between the two shapes, i.e. the ground truth (x_{gt}) and the optimized shape (x). Two distance measures defined over landmarks are used to obtain the convergence performance. The first one is called point to point error, defined as the Euclidean distance between each corresponding landmark:

$$D_{pt.pt.} = \sum \sqrt{\left(x_i - x_{gt,i}\right)^2 + \left(y_i - y_{gt,i}\right)^2} \quad (1)$$

The other distance measure is called point to curve error, defined as the Euclidean distance between a landmark of the fitted shape (x) to the closest point on the border given as the linear spline, $r(t) = \left(r_x(t), r_y(t)\right), t \in [0,1]$, of the landmarks from the ground truth (x_{gt}):

$$D_{pt.crv.} = \frac{1}{n}\sum_{i=1}^{n} \min_t \sqrt{\left(x_i - r_x(t)\right)^2 + \left(y_i - r_y(t)\right)^2} \quad (2)$$

We have calculated these errors for all for datasets (from Set 1 to Set 4). The AAM searching is known to be very sensitive to the selection of initial configuration. We tested the proposed method against the selection of initial configuration. We translate, rotate and scale initial configurations and see how the proposed method can handle the poor initialization. We made 10 experiments for each test image with different initializations and took the average error. These experiments include mean-shape configuration, ±5 degrees rotation, scaling by 0.85 and 0.95, translation by 10% in x and y directions. Table.1 summarizes the averages of point-to-point and point-to-curve errors when classical AAM search is used without any illumination normalization. Point-to-point and point-to-curve errors obtained by the proposed illumination normalization method are much less than the errors obtained by the classical AAM (Table.2).

Ratio-image method is not suitable for AAM searching, at least for the first iterations of the algorithm. Let's suppose that we start searching in a position far away from the ground truth location. The model synthesizes a face that best fits the current location. Then the textures of the synthesized face and corresponding part in the test image are analyzed and an error coefficient is computed, reflecting the similarity degree of the two textures. We normalize the corresponding texture in the test image before computing the error. The main problem with the ratio-image method is that when it is applied to a region of an image that is not face-like, the normalization result will have a lot of information of the mean-face, putting in other words it will be mean-face-like. Thus the error will be much smaller than the real one, and it will introduce false alarm in the searching process creating additional local minima. On the other hand, the histogram based normalization method will never change the general aspect of an image, only the pixel intensities follow a different distribution. Thus the chances of introducing false alarms are reduced using this normalization method. The ratio-image can produce very good results provided that the shape is already aligned. But this is not the case in AAM searching. We assume that the best fit returned by the searching algorithm using histogram-based normalization is a good approximation of the real face, and thus the alignment requirement is satisfied.

Table.1 Standard AAM fitting performance.

	Set 1	Set 2	Set 3	Set 4
Pt.pt.	4.9±0.20	11.4±0.57	19.4±0.58	36.6±1.64
Pt.Crv.	2.9±0.11	6.8±0.33	12.9±0.36	33.2±1.44

Table.2 Proposed AAM fitting performance.

	Set 1	Set 2	Set 3	Set 4
Pt.pt.	4.1±0.12	8.06±0.34	13.03±0.41	21.3±0.58
Pt.Crv.	2.4±0.08	5.24±0.23	8.76±0.29	14.71±0.42

(a) (b) (c) (d)

Figure.5 Searching results First row is the classical AAM searching results, second row is the proposed method (a) Initial configuration (b) Mean face (c) Searching result obtained in the 3th iteration (d) Searching result obtained in the 6th iteration

3. Pose Normalization

Pose normalization is required before recognition in order to reach acceptable recognition rates. There are several works related to pose normalization. Blanz and Vettel [11] use a statistical 3D morphable model to tackle with pose and illumination variations. Since their method requires textured 3D scans of heads, it is computationally expensive. Cootes et al constructed three AAMs which are called as View-based AAMs [9]. We developed AAM based pose normalization method which uses only one AAM. There are two important contributions over the previous studies. By using the proposed method:

i. One can synthetically generate appearances for different poses when only frontal face image is available.
ii. One can generate frontal appearance of the face when there is only non-frontal face image is available.

Next section explains the proposed pose normalization and generation method.

3.1. Pose Generation from 2D Images

The same variation in pose imposes similar effect on the face appearance for all individuals. Deformation mostly occurs on the shape whereas the texture is almost constant. Since the number of landmarks in AAM is constant, the wireframe triangles are translated or scaled as pose changes. So as we change pose, only wireframe triangles undergo affine transformation but the gray level distribution within these triangles remains the same. One can easily generate frontal face appearance if AAM is correctly fitted to any given non-frontal face of the same individual provided that there is no self-occlusion on face. Self-occlusion usually is not a problem for angles less than *±45*.

For 2D pose generation, we first compute how each landmark point translates and scales with respect to the corresponding frontal counterpart landmark point for 8 different poses, and obtain a ratio vector for each pose. We use the ratio vector to create the same pose variation over the shape of another individual. Appearances are also obtained through AAM using synthetically generated landmarks. These are shown in Fig.6. First column in Fig.6 shows the frontal faces and the second column shows appearances for various poses. It is important to note that the generated faces contain no information about the individual used in building the ratio matrix.

3.2. Training AAM for Pose Normalization

An AAM model trained by using only frontal faces can only fit into frontal faces well and fail to fit into non-frontal faces. Our purpose here is to enrich the training database by inserting synthetically generated faces at different poses so that AAM model trained by frontal faces can now converge to images at any pose.

We manually labeled 73 landmarks on 4920 images. Let us denote the landmark points on i^{th} frontal image as $S_i^0 = \left(\left(x_{i,1}, y_{i,1}\right), \left(x_{i,2}, y_{i,2}\right), \ldots, \left(x_{i,K}, y_{i,K}\right)\right) \in R^{2K}$ where $i = 1, 2, \ldots, N$. N is 4920 and K=73 in our database. The shape-ratio vector explained in the previous subsection (3.1) is defined between the *p*-posed shape and the frontal shape as

$$r_p(S^p, S^0) = \left(\left(\frac{x_{p,1}}{x_{0,1}}, \frac{y_{p,1}}{y_{0,1}}\right), \ldots, \left(\frac{x_{p,K}}{x_{0,K}}, \frac{y_{p,K}}{y_{0,K}}\right)\right)$$

Shape of any unseen individual at pose *p* can now be easily obtained from frontal shape using shape-ratio vector r_p as

$$\hat{S}^p_{unseen} = r_p S^0_{unseen}.$$

We synthesize shapes from frontal-view images in the database for P=8 different poses as,

$$\hat{S}_i^p = r_p S_i^0, i=1,2,\ldots,10, \text{ and } p=1,2,..,8.$$

AAM shape component is constructed from these aggregated shapes $\hat{S}_i^p$ and S_i^0 by applying principal component analysis as $S = \bar{S} + Q_s s$ where $\bar{S}$ is the mean shape, Q_s contains k eigenvector of the covariance matrix corresponding to the highest k eigenvalues.

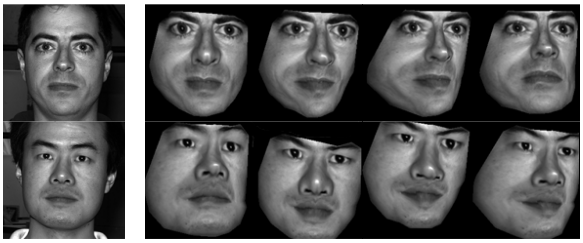

(a) (b)

Figure.6 Synthetic pose generation from frontal face:
a) Frontal face, b) Synthetically generated non-frontal faces.

Next step is to warp each face in the training database to mean shape ($\overline{S}$) and apply principal component analysis to the texture this time as $T = \overline{T} + Q_t t$ where $\overline{T}$ is called mean face. Any shape (*S*) and texture (*T*) can be steadily mapped to the AAM subspace as $s = Q_s^T \left(S - \overline{S} \right)$ and $t = Q_t^T \left(T - \overline{T} \right)$.

AAM is comprised of both shape (Q_s) and texture (Q_t) subspaces. Any change in face shape leads to a change in face texture and vice versa. Face appearance (*A*) is dependent on shape and textures. This dependency is expressed as $A = \left[\Lambda s \quad t \right]^T$. In order to exploit the dependency between shape and texture modeled by the diagonal matrix (Λ), one further PCA is applied to the shape and texture components collectively and we obtained the combined model called appearance model as $A = Q_a a$. Any appearance is obtained by a simple multiplication as $a = Q_a^T A$.

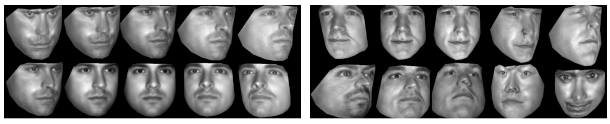

Figure.7 Randomly synthesized faces from leading 5 AAM parameters.

In order to show how rich representation AAM provides us, we used the first 5 coefficients and select random points in 5-dimensional space. The corresponding faces are plotted in Fig.7. Even this simple experiment proves that AAM trained as explained above can generate pose variations not governed by any shape ratio vector (r_p). We also conducted another experiment to see how close we fit into unseen faces at different poses. Fig.8 summarizes the alignment results for these unseen faces.

4. Experimental Results

We also analyze how the proposed alignment method affects the recognition performance. We used the following feature spaces in our experiments: PCA, LDA. Randomly selected 25 images of each person from Set 1 dataset are used in training. All datasets (Set 1 through Set 4) contain faces of all poses. The remaining faces in Set 1 dataset are used as test data. Recognition rates for two feature spaces (i.e., PCA and LDA) in Set 1-4 are plotted in Fig.9 for increasing dimensions. The recognition rates obtained when the original images are used as input to the classifier are denoted as ORG-PCA and ORG-LDA. The recognition rates obtained when the images restored by RI are used as input are denoted as RI-PCA and RI-LDA. Finally, the recognition rates obtained when the images restored by HF are used as input are denoted as HF-PCA and HF-LDA. PCA is known to be very sensitive to misalignment in faces. Our experimental studies also verify this behavior. When the original images are used, the PCA recognition rates for all sets are poor. LDA is more successful for dimensions closer to 9. ORG-PCA reaches to 74.36% at most, while ORG-LDA reaches to 91.26% at most for Set 1. This performance drops to 30.99% for ORG-PCA and to 41.13% for ORG-LDA for Set 4.

One important observation is that AAM alignment with histogram fitting always leads to better recognition rates in all test sets (Set 1- 4) compared to case where original faces are used and ratio-image normalization is used right after the AAM alignment. Another advantage of the proposed method is that similar recognition performance is obtained at lower dimensions. Recognition rate for ORG-LDA is just 32.81% while LDA performance for the proposed approach (called HF-LDA) is 83.38% when the dimension is set to 3. ORG-LDA catches this rate when the dimension is set to 5.

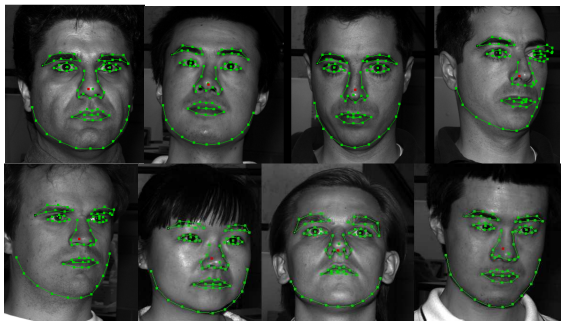

Figure.8 Face alignment result for unseen faces.

For the challenging test set, i.e. Set 4, both ORG-LDA and ORG-PCA fails. The recognition rate is at most 30.99% for ORG-PCA and 41.13% for ORG-LDA. On the other hand, HF-PCA reaches at most to 76.20% and HF-LDA reaches at most to 82.68%. This is a significant improvement when compared to the results obtained without applying any preprocessing (41%). Note that all test sets include faces of 8 different poses selected from Yale B dataset.

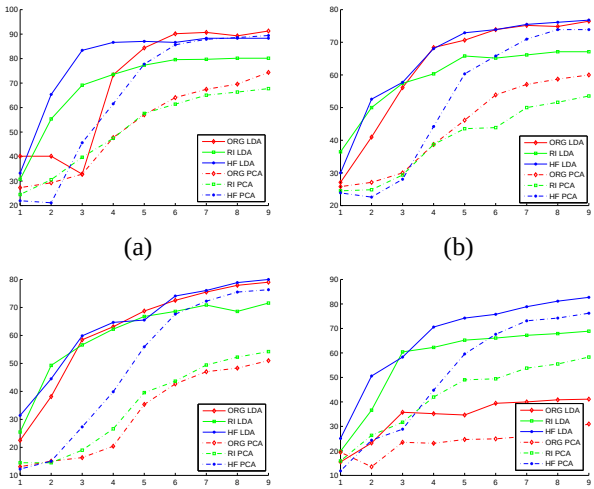

Figure.9 PCA and LDA recognition rates for Set 1 (a), Set 2 (b), Set 3 (c), and Set 4 (d) when original face (ORG), Ratio Image (RI) and the proposed restoration (HF) are used.

Figure.10 Initialization (first row) and alignment/restoration results of proposed method (second row) for different pose and illumination variations.

5. Conclusion

In this study we developed AAM based face alignment method which handles illumination and pose variations. The classical AAM fails to model the appearances of the same identity under different illuminations and poses. We solved this problem by inserting histogram fitting based normalization into the searching mechanism and inserting different poses of the same identity into the training set.

From the experimental results, we showed that the proposed face restoration scheme for AAM provides higher accuracy for face alignment in point-to-point error sense. Recognition results based on PCA and LDA feature spaces showed that the proposed illumination and pose normalization outperforms standard AAM.

6. Acknowledgement

This work was partly supported by the National Scientific and Research Council of Turkey; project no: EEEAG-104E121 and the State Planning Agency (DPT) of Turkey. We would like to thank Florin S. Telcean for his contributions and suggestions.

7. References

[1] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active Shape Models-their training and application," Computer Vision and Image Understanding, 61(1), pp. 38-59, 1995.

[2] T.F. Cootes, G. Edwards, and C.J. Taylor, "Active appearance models," IEEE Trans. on PAMI, vol. 23, no.6, pp. 681-685, 2001.

[3] Stegmann M. B., Ersboll B. K., Larsen R., FAME - A Flexible Appearance Modeling Environment, IEEE Trans. on Medical Imaging, vol. 22(10), pp.1319-1331, 2003

[4] W. Zhao and R. Chellappa, "SFS Based View Synthesis for Robust Face Recognition," Proc. 4th Conf. on Automatic Face and Gesture Recognition, 2000.

[5] H. Wang, Z. S. Li, Y. Wang, W. Zhang, "Illumination Modeling and Normalization for Face Recognition", In Proc. AMFG 2003, pp.104-111, 2003.

[6] R. Basri and D. Jacobs, "Photometric Stereo with General, Unknown Lighting", CVPR, vol.2, pp.374-381, 2001.

[7] A. Shashua and T. Riklin-Raviv, "The Quotient Image: Class-Based Re-Rendering and Recognition With Varying Illuminations", IEEE Trans. on PAMI, pp. 129-139, 2001.

[8] W. Zhao and R. Chellappa, "Face Processing: Advanced Modeling and Methods", Academic Press, Elsevier, 2006.

[9] T. Cootes, G. Wheeler, Walker K., and Taylor C., "View based active appearance models", Image and Vision Computing, 20:657–664, 2002.

[10] X. Hou, S. Li, H. Zhang, and Q. Cheng, "Direct appearance models", In CVPR, pp. 828–833, 2001.

[11] V. Blanz, T. Vetter, "A morphable model for the synthesis of 3D faces", 26th Conf. on Computer graphics and interactive techniques, p.187-194, 1999.

[12] D. H. Liu, K. M. Lam, L. S. Shen, "Illumination invariant face recognition", Pattern Recognition, 38(10):1705-1716, 2005.

[13] T. Jebara, "3D Pose Estimation and Normalization for Face Recognition", B. Thesis, McGill Centre for Intelligent Machines, 1996.

[14] Yale B Web Site: http://cvc.yale.edu/projects/yalefacesB/yalefacesB.html.