3D Face Recognition Using 3D Alignment for PCA

Trina Russ
Security Technology Dept.
Sandia National Labs[1]
Albuquerque, NM
tdruss@sandia.gov

Chris Boehnen
Computer Science and Engineering
University of Notre Dame
Notre Dame, IN
cboehnen@nd.edu

Tanya Peters
Security Technology Dept.
Sandia National Labs[1]
Albuquerque, NM
thpeter@sandia.gov

Abstract

This paper presents a 3D approach for recognizing faces based on Principal Component Analysis (PCA). The approach addresses the issue of proper 3D face alignment required by PCA for maximum data compression and good generalization performance for new untrained faces. This issue has traditionally been addressed by 2D data normalization, a step that eliminates 3D object size information important for the recognition process. We achieve correspondence of facial points by registering a 3D face to a scaled generic 3D reference face and subsequently perform a surface normal search algorithm. 3D scaling of the generic reference face is performed to enable better alignment of facial points while preserving important 3D size information in the input face. The benefits of this approach for 3D face recognition and dimensionality reduction have been demonstrated on components of the Face Recognition Grand Challenge (FRGC) database versions 1 and 2.

1. Introduction

Facial recognition (FR) is a suitable biometric for many applications because of its non-intrusive nature, ease of integration into existing systems, societal acceptance, and potential to identify individuals at a distance without subject cooperation. Most of the current efforts in FR use a 2D representation of the face [12], however, these systems are generally not robust to variations in facial pose, scene lighting, and facial expression, requiring images to be taken in a controlled environment [9]. To address these issues researchers are examining the benefits of 3D data for the FR task [3]. 3D FR systems are believed to be less sensitive to lighting and facial pose variations. In addition, the geometric information available is expected to provide more discriminatory features for FR, including size information which is lacking in most 2D approaches.

Principal component analysis (PCA) has been extensively used for the analysis, compression, modeling, and recognition of objects because of its simplicity and ability for dimensionality reduction. In particular, it has been used for facial recognition [2, 4, 11, 14, 18, 20]. The goal of PCA is to derive an optimal basis from a set of training examples (representing signals in class **C**) such that training examples can be represented as a linear combination of basis vectors. For recognition problems, it is also desirable for the derived basis to have good generalization performance. For instance, the basis should also accurately represent new signals (not in the training set) in signal class **C** as a linear combination. To achieve both data accuracy and efficiency in representing new within class signals, both sufficient variability in the training set and proper correspondence of signal features (in concert with linear combination principals) are needed.

In 2D face analysis this is obtained by a normalization procedure which typically requires certain attributes to be constant across faces. For instance, for a 2D face one often restricts the distance between eyes to be fixed and the nose tip to be placed at a particular pixel in the 2D image [13]. For 3D face recognition, PCA has been applied to range images of faces [4, 11, 18, 20]. Range images are often noted as 2.5D representations and two issues may arise when using them for 3D PCA recognition: 1) if normalization is performed using typical 2D methodology, size information which is important for the recognition process is eliminated, 2) if normalization is not performed then proper alignment of facial points necessary for good generalization is not achieved. Accurate representation of new untrained signals in the signal class will typically require more model parameters and/or training examples. We

[1]Sandia is a multi-program laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

address these fundamental issues with a truly 3D approach for PCA face recognition.

The presented approach enables a good alignment of 3D points across faces while still preserving face size information important for the recognition process. This is achieved using a generic reference face model that can be scaled in 3D to best fit an input face. By scaling the model, opposed to the input face, proper size information of the face of interest can be encoded and better alignment of that face to the reference can be performed. Essentially, the reference face is the basis for a new coordinate system that provides a mapping between distinct faces to be encoded in the PCA decomposition. Input face to reference face alignment is achieved using the Iterative Closest Point (ICP) algorithm, and correspondences are obtained using a special normal search criterion. The utility of this approach for 3D face recognition has been demonstrated on components of the Facial Recognition Grand Challenge (FRGC) datasets versions 1 and 2. Results demonstrate both good generalization and data dimensionality reduction.

2. Related Work

A comprehensive review of 3D FR research is presented in [3]. Therefore, we will not perform a thorough review, but mention those findings closely related to our 3D approach for face recognition.

Chang *et al.* [4] examined the benefits of using 3D and 2D+3D approaches over 2D using a PCA algorithm. Results were presented using a database very similar to the FRGC v1.0 database. For a single probe database, they achieved a rank one recognition rate of 89%, 94.5%, and 98.5% for 2D, 3D, and 2D+3D, respectively. In [5], Chang *et al.* presented results of PCA and ICP algorithms on FRGC v2 containing over 450 subjects for comparison to a new approach more robust to variations in facial expression. For neutral expressions, average rank one identification scores were 77.7% and 91% for PCA and ICP, respectively. Analysis of non-neutral face expressions resulted in average rank one identification rates of 61.3% and 61.5% for PCA and ICP, respectively.

In [2], Vetter and Blanz utilized a 3D model based on PCA to address the problem of pose variation for 2D face recognition systems. A 3D model was generated from real 3D scans. To obtain a fixed ordering of points for PCA alignment, a reference face was used. Correspondence to the reference frame was obtained using an optical flow based algorithm that examined 2D texture and 3D shape information. 3D PCA model parameters initialized by features extracted from a 2D image served as the basis for recognition.

Other approaches to PCA face recognition using range images were presented in [11, 18, 20]. These approaches corrected for orientation and preserved 3D size information similar to [4]. Using a PCA basis derived from one or more examples of subjects in the test set, [11] obtained a verification rate of 99.17% with a 0 false accept rate (FAR) analyzing 50 subjects, and [20] obtained a rank one identification rate of 90% analyzing 37 subjects. In [18], independent training and test sets were used to obtain an equal error rate (EER) of 12.7% analyzing 100 subjects with variations in expression and pose.

Iterative matching algorithms for 3D face matching using range images were presented in [7] and [10]. An EER of 3.5-4.4% was obtained in [7] analyzing 18 subjects and a verification rate of 93.5% at a FAR of .001 analyzing 198 subjects was achieved in [10]. Deformations of faces from an annotated face model were analyzed for FRGC v2 in [19], achieving verification rates of 94.45% and 60.69% at .001 for neutral and non-neutral expressions respectively.

Our approach seeks to move away from range-based algorithms [4, 7, 11, 18, 20] to a full 3D model. This approach allows proper consideration of all coordinates in a 3D space. In addition, avoidance of an image grid enables additional flexibility in how the actual points on the face are aligned. We are distinguished from [2] in that our focus is not 2D face recognition, but how to perform accurate 3D face recognition. While [2] utilizes 2D data to synthesize an approximate 3D model, we start with 3D data. Because of this we choose to avoid optical flow based correspondence to a reference, and instead, perform 3D point registration to a 3D reference face.

3. 3D Face Modeling

3.1. Principal Component Analysis

Principal component analysis is a common method for deriving an optimal set of orthogonal basis functions from a linear class of signals [15]. This decomposition enables the signal class to be represented as a linear combination of the PCA basis vectors. We apply this technique to 3D facial analysis. Given M sets of 3D coordinates $\mathbf{S} = \{\mathbf{s}_1, \mathbf{s}_2, ..., \mathbf{s}_M\}$ representing M different faces, we seek to derive a 3D face basis $\mathbf{V}$ to represent signals in $\mathbf{S}$ as a linear combination. In order to achieve this, the coordinates in each dataset should be in correspondence and the same number of points should reside in each set $\mathbf{s}_i$. This modified dataset (discussed in section 3.2) is referred to as $\mathbf{X}$ and is encoded in the PCA decomposition. Thus, any dataset $\mathbf{x}$ can be

reconstructed from the PCA basis by (1) to produce $\hat{\mathbf{x}}$. The reconstructed signal $\hat{\mathbf{x}}$ is optimal in the sense that the L2 norm between $\mathbf{x}$ and $\hat{\mathbf{x}}$ is minimized as shown in (2).

$$\hat{\mathbf{x}} = \bar{\mathbf{x}} + \sum_{i=i}^{M-1} \alpha_i \mathbf{v_i} \tag{1}$$

$$\left\| \hat{\mathbf{x}} - \bar{\mathbf{x}} - \sum_{i=i}^{M-1} \alpha_i \mathbf{v_i} \right\| \tag{2}$$

The PCA basis $\mathbf{V}$ is determined as follows. The mean face $\bar{\mathbf{x}}$ (3) is subtracted from $\mathbf{x_i}$ to produce $\mathbf{a_i}$ (4). These vectors define a matrix as shown in (5). The PCA basis vectors can be determined by computing the eigenvectors of the covariance matrix $\mathbf{C}$ (6), obtained using the Singular Value Decomposition of $\mathbf{A}$.

$$\bar{\mathbf{x}} = \sum_{i=i}^{M} \mathbf{x_i} \tag{3}$$

$$\mathbf{a_i} = \mathbf{x_i} - \bar{\mathbf{x}} \tag{4}$$

$$\mathbf{A} = \{\mathbf{a_1}, \mathbf{a_2}, ..., \mathbf{a_M}\} \tag{5}$$

$$\mathbf{C} = \frac{1}{m} \mathbf{A}\mathbf{A}^T \tag{6}$$

Given an input face $\mathbf{x}$ the PCA model coefficients can be determined by solving the least squares problem in (2). Equation (2) can be rewritten as $\mathbf{a} = \mathbf{V}\boldsymbol{\alpha}$ and solved by computing the pseudo-inverse of $\mathbf{V}$ (7). The utility of these coefficients $\boldsymbol{\alpha}$ is examined for facial recognition in section 4.2.

$$\boldsymbol{\alpha} = \mathbf{V}^{+}\mathbf{x} \tag{7}$$

3.2. Correspondence Calculation

The correspondence calculation is a critical component of PCA and our approach to this is one of the unique contributions of this work. Ideally, like features should be corresponded (i.e. eyes are aligned with eyes, nose with nose, etc.). However, in practice this is difficult and good performance can be achieved with a coarse approximation to this. For example, 2D techniques seek alignment of nose positions; however, other features such as the mouth and chin may not be exactly aligned. Despite this, PCA is able to produce good results for 2D recognition. Similarly, we do not seek exact alignment of prominent facial features, but a good approximation to it over similar face geometries.

3.2.1. Reference Face Utilization

The ordering of points for PCA is determined by alignment of an input face to a generic reference face as adopted in [2]. By ordering all of the points according to points on a reference face, all points are in a common consistent coordinate system from which PCA model parameters can be derived. An attractive attribute of this approach is that in a matching paradigm when one is searching for an input face in a gallery database, the input face does not have to be aligned with each face in the gallery for matching to occur, enabling fast database searches.

To enable correspondence of new faces to the reference face, the reference face is scaled with respect to an input face. The five features shown in Figure 1 provide a basis for determining scaling parameters. An automated technique to detect these features has been developed in [6] based on radial symmetry. Scales in the x, y, and z dimensions can be determined by calculating the ratio of feature distances between the reference face and input face. For the experiments conducted in Section 4, scaling was only performed along the vertical direction with the nose location at the origin. Using the inner eye corners and nose to scale in the x and z directions significantly impacted recognition performance. We believe this is because the inner eye corners are closely located to each other; thus, small errors may have a significant impact on the scaling parameter resulting in incorrect scaling of the reference. In order to assess our approach without feature detection errors, ground truth locations were used for the experiments in this paper.

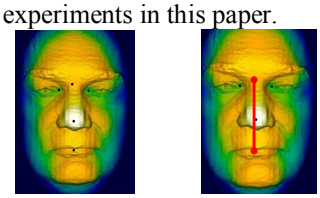

features scale parameter

Figure 1. Example face features used to determine reference face scaling prior to alignment.

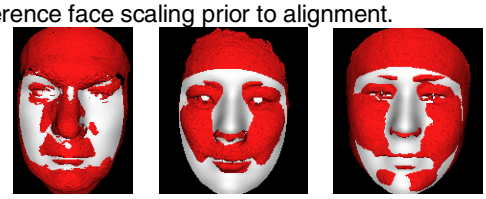

Figure 2. Reference face (white) to input face (red) registration.

Reference to 3D face alignment is achieved using the Iterative Closest Point [1] to minimize the root mean-squared error between the two representations. Examples of the reference face to input face alignment are shown in Figure 2. Although, the faces are of varying sizes and shapes, the reference face is still well aligned with each input face.

3.2.2. Normal Search Correspondence

After face alignment, a corresponding point on the input surface is found for each point on the reference

surface by intersecting a given reference point's normal with the input surface. This tends to give improved performance over a nearest neighbor search because it enables an improved sampling of the input face's surface geometry as can be seen in Figure 3. However, for a normal search algorithm to provide good results, the reference face must have a low noise level amongst the surface normals. This is obtained by applying smoothing and other processing to the reference face which can be generated offline. An example of the result of the correspondence stage can be viewed in Figure 4. The same ICP algorithm is performed for both examples; only the criteria for selecting the appropriate match to a reference point changes. The correspondence result in Figure 4a selected the nearest neighbor, and the result in Figure 4b selected the point that intersects the surface in the normal direction. Notice the improved definition around the mouth and nose for the normal search method. Also, notice how the nearest neighbor search distorts the upper lip in Figure 4b. Added definition and accuracy at this stage will benefit the final recognition stage.

3D face representations often contain erroneous data points because of limitations in 3D sensor technology. To address this problem, smoothing and/or processing is performed prior to ICP alignment and after 3D point correspondence. The 3D points are triangulated across holes to provide a valid surface for the reference face's normal to intersect during correspondence. In cases where the normal search correspondence strays too far from the closest point on the surface, the nearest neighbor is used. To correct erroneous correspondences, the post-processing step utilizes neighboring correspondences.

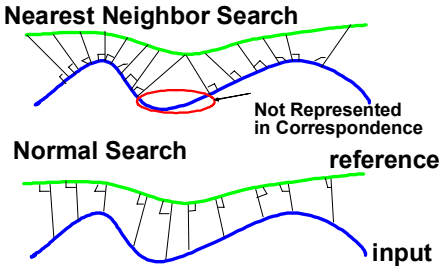

Figure 3. An example of the nearest neighbor and normal search correspondence methods.

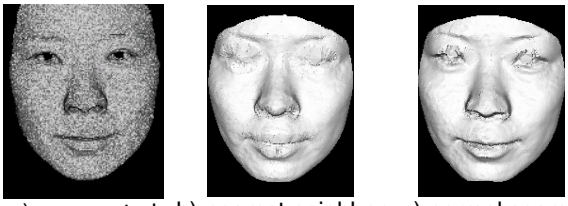

a) segmented b) nearest neighbor c) normal search

Figure 4. Examples of reference based correspondence reconstructed datasets using different correspondence criteria.

3.2.3. Reference Face Generation

The reference face can be obtained in a number of different ways. For our experiments, a face from the FRGC database was processed in Raindrop Geomagic 6. The choice of the selected face for processing was random; a number of faces would probably work equally well. To derive a reference face, processing was performed both globally and regionally including: surface decimation, smoothing, and relaxation.

At this stage, the goal was to create a smooth reference model containing the defining characteristics of the face (i.e. eyes, nose, mouth area, etc.); however, we did not want the facial features to be too distinct. The use of an abstracted smooth reference face allows the ICP process to better align the reference face to a variety of different face geometries and aids in normal search correspondence. In addition, parts of the face were cropped to remove and avoid the inclusion of undesirable areas such as hair. After extensive processing, the reference face contained approximately 23,000 points. A picture of the FRGC face, before and after processing, is shown in Figure 5.

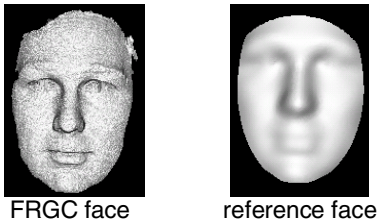

FRGC face reference face

Figure 5. An example of a 3D face and a derived reference face used for correspondence.

4. Results

Analysis of the presented system is performed for 3D facial verification and identification on multiple probe databases from the FRGC database versions 1 and 2 [8]. FRGC v1 consists of images of 200 subjects acquired in the spring of 2003. Of the 200 subjects, 198 participated in more than one session allowing the first image to be used as a gallery image (enrollment), and subsequent acquisitions to be used as probe images (test image against the gallery). For this database, 745

probe images were analyzed and hereafter referred to as the spring 2003 database. The FRGC v2 includes FRGC v1 and images acquired during the fall of 2003 and spring of 2004. FRGC v2 contains significant variations in facial expression. Only images of subject with more than 4 images acquired in fall 2003 were used in the following analysis. This resulted in 242 neutral gallery images and 1287 probe images with neutral and non-neutral expressions. Hereafter, this database is referred to as fall 2003.

4.1. 3D Face Synthesis

3D PCA face synthesis is the process of generating a new face based on the model; however, this process can be impacted by the quality of the segmentation, the ability to properly generate good correspondences to the reference model, and the set of training faces from which the PCA face basis is derived. In this section, the face generated by correspondence to our reference face is called the reconstructed face. In general, the best PCA can achieve is the accuracy of the reconstructed face. However, if the PCA training images are free of correspondence errors then the model can potentially overcome inaccuracies in the reconstructed face introduced by the correspondence stage. Face segmentation was performed by cutting a sphere near a manually marked nose location and performing surface clustering to remove extraneous unconnected regions that fall within this area. Subsequently, the correspondence procedure described in section 3.2 was performed and PCA model coefficients were calculated.

Figure 6 shows examples of reconstructed and synthesized faces generated with a PCA training set of approximately 200 faces representing unique subjects. We demonstrate the utility of the approach for three cases: a face image included in the PCA training set, a face image whose subject was included in the training set (i.e. different data set of subject), and the face image of a subject not included in the training set. For all three cases, the quality of the reconstructed and synthesized faces is very good; although, in practice correspondence errors may occur.

4.2. Facial Recognition Analysis

In this section, we examine the utility of the PCA model parameters for face recognition by examining both receiving operating characteristic (ROC) curves and cumulative match curves (CMC's) [9]. The ROC curve graphs the probability of correct verification versus FAR and provides an indication of discrimination ability as various thresholds are used to define correct verification. The CMC displays the probability of identification if the top K match results were returned. This provides an indication of how close one is to obtaining the correct match if the rank one match was incorrect. In a real world system, the top K matches can be given to security personnel to make the final identification decision.

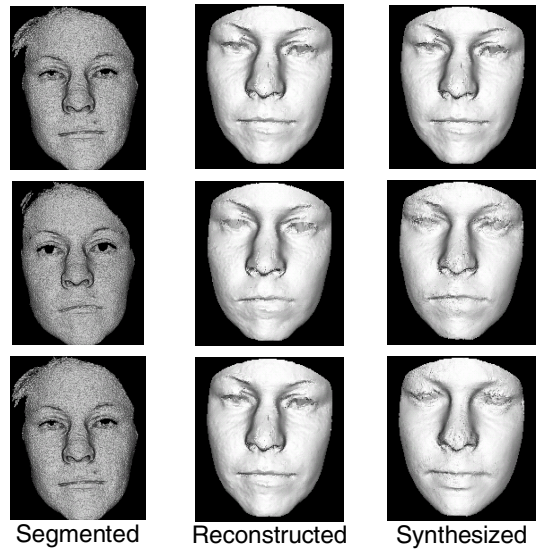

Figure 6. 3D PCA Synthesis Examples for (Top) dataset in the training set (Middle) dataset not in training set, but subject in the training set (Bottom) dataset and subject not in training set

The match metric used for facial recognition is one minus the cosine of the angles between the PCA model parameters; however, whitening is not performed as in [13]. Given a vector u representing the model parameters of face A and v representing the model parameters of face B, this metric is defined by $Q(u,v)$ in (8). When the vectors are identical the metric is 0.0 and when they are orthogonal, it is 1.0. In practice, this metric performed better than the L2 norm.

$$Q(u,v) = 1 - \frac{\sum_{i=i}^{M-1} u_i v_i}{\sqrt{\sum_{i=i}^{M-1} u_i^2}\sqrt{\sum_{i=i}^{M-1} v_i^2}} \tag{8}$$

PCA is known for its efficiency in dimensionality reduction. Thus, we examined FR performance using k% of the most important parameters in matching. FR performance is analyzed on the neutral images from the spring and fall 2003 databases containing 198 and 202 gallery images and 765 and 565 probe images, respectively. The gallery sets from each database were used for PCA training. Figure 7 and Figure 8 show the ROC's for both the spring and fall 2003 databases. Both ROC curves demonstrate verification rates of approximately 95% for a FAR of .001. In addition, the rank one identification rate varies from 91-97% for the spring and fall 2003 databases (Figure 9 and Figure 10). Notice little degradation in performance when

using 100% to 20% of the basis coefficients (~200 coefficients).

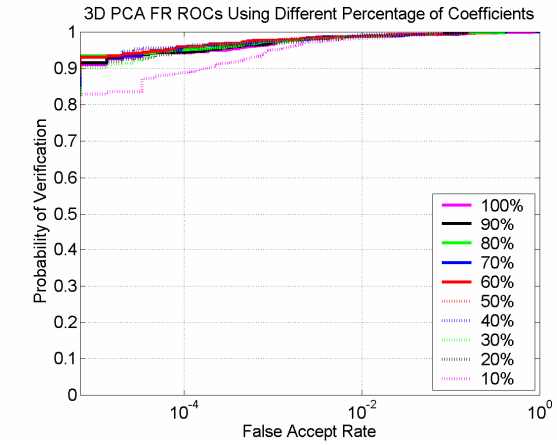

Figure 7. Spring 2003 neutral ROC curves

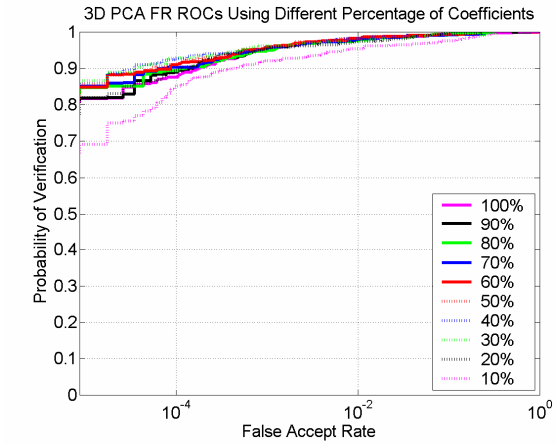

Figure 8. Fall 2003 neutral ROC curves

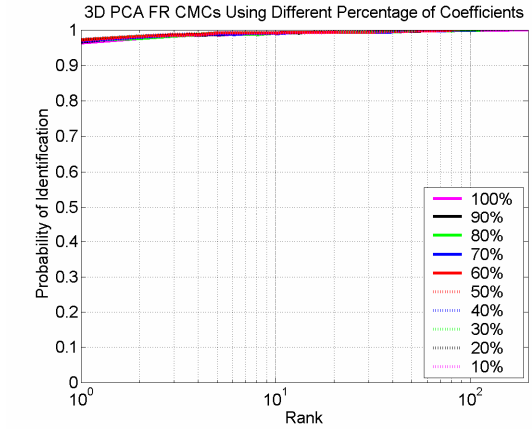

Figure 9. Spring 2003 neutral CMC

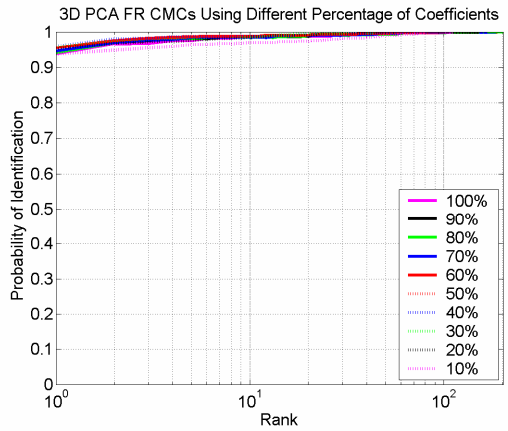

Figure 10. Fall 2003 Neutral CMC

4.3. 3D PCA Generalization Analysis

A benefit of the proposed approach over range-based PCA modeling algorithms is the potential for greater accuracy when provided with new untrained data. In order to test this capability, we varied the training set size for a combined spring and fall 2003 database. This resulted in a total of 310 neutral gallery subjects and 1450 neutral probe images. The training set was obtained by taking the first N images of the gallery set where N varied between 50 and 310. A synthesized example showing the impact of training set size on face reconstruction is shown in Figure 11. Note this subject was not included in any of the training sets except for when N=310.

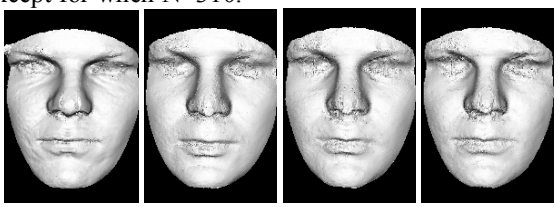

Reconstructed N=50 N=100 N=200

Figure 11. PCA synthetic reconstruction using different training set sizes

Figures 12 and 13 show the ROC curves and CMCs generated for the combined spring and fall 2003 neutral databases. The graphs shown are only using 60% of the N-1 model coefficients. With this approach good performance is obtained using a small but representative training set achieving verification rates of 92-96% at a .001 FAR and rank one identification rates between 94-95%. This is an improvement over the range-based PCA results presented in [4] and [5]. While good results were demonstrated in [4], performance dropped significantly in [5] when considering additional images and subjects. Rank one recognition rates ranged from 73.3% to 83.5% analyzing neutral expressions (~ 449 subjects). We have analyzed a subset of this data, but contend that

our results are significant because of the amount of improvement achieved and the consistency of the results over variations in training set size and the number of coefficients used.

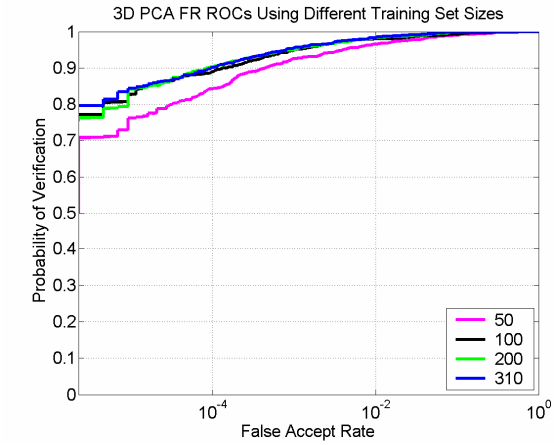

Figure 12. Fall and spring 2003 ROC curves using different training set sizes

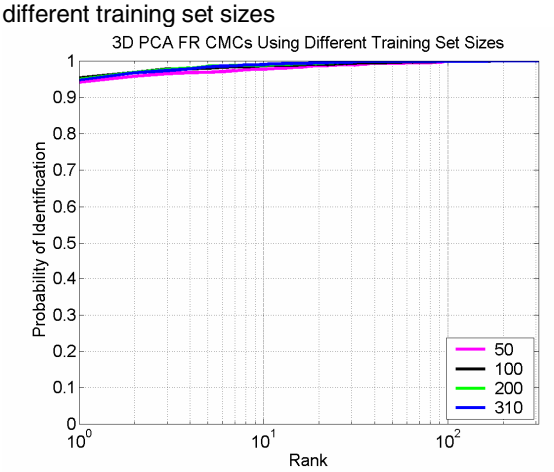

Figure 13. Fall and spring 2003 CMC using different training set sizes

4.4. Facial Expression Analysis

Dealing with facial expression variations is a challenge for many 2D and 3D recognition systems because they cause non-rigid deformations in the facial surface, but it is worth noting the utility of the proposed method in addressing this problem.

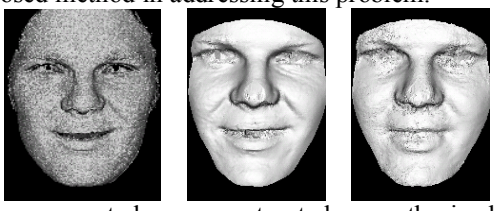

segmented reconstructed synthesized

Figure 14. An example of PCA synthesis of a non-neutral expression using a basis trained on mostly neutral expression images.

To examine this issue, a gallery of neutral images is matched to a probe set of non-neutral images from the fall 2003 database. The fall and spring 2003 databases were combined to extract a set of 310 unique subjects used to train PCA. The probe set contains 840 images of which only 753 have a match in the fall 2003 neutral gallery of 202 subjects. An example of PCA synthesis of a non-neutral expression with a mostly neutral expression training set is shown in Figure 14. Note the similarity in facial structure, although the amount of smiling between the original and synthesized versions is noticeably different.

The ROC curve and CMC for this experiment are shown in Figures 15 and 16. There is a considerable decrease in performance when dealing with non-neutral expressions as expected; however, a rank one identification rate of 82.6% and a verification rate of 91% with a .001 FAR is achievable using 60% of the coefficients. Furthermore, the 3D model has potential to correct for variations in facial expression similar to [17].

5. Conclusion and Future Work

We have presented a general approach for 3D face recognition based on PCA. It avoids elimination of size information by scaling a 3D reference to enable alignment of facial points important for PCA training, synthesis, and recognition. This work represents an improvement upon the work presented in [4, 5] in recognition accuracy, dimensionality reduction, and generalization to new faces. The presented approach also has benefits over the iterative registration algorithms presented in [1, 5, 7, 10] during a 1 to N match in that it only requires the registration of the probe to a reference, speeding up the algorithm considerably. In future work, we intend to examine the FRGC v2 database in its entirety (i.e. spring 2004) and use the 3D model as a foundation for dealing with the problem of facial expression variation [17]. We also seek to examine robust scaling parameters to be applied across other dimensions of the face besides the vertical one and examine the impact of scaling errors or recognition accuracy. Also, we intend to replace the ICP reference alignment algorithm with the 2D range Hausdorff approach described in [10], which can be computed in O(N) opposed on $O(N^2)$.

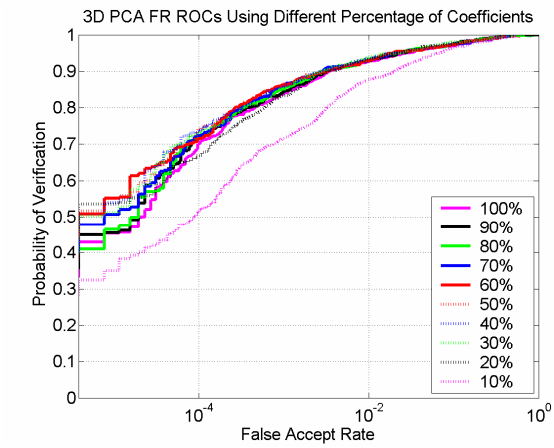

Figure 15. Fall 2003 non-neutral ROC

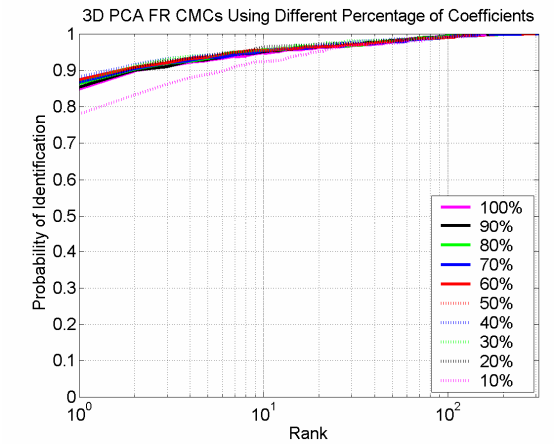

Figure 16. Fall 2003 non-neutral CMC

References

[1] P. Besl and N. McKay. A method for registration of 3d shapes. *IEEE PAMI*, 14(2): 239-256, February 1992.

[2] V. Blanz and T. Vetter. Face recognition based on fitting a 3d morphable model. *IEEE PAMI*, 25(9):1063-1074, September 2003.

[3] K. W. Bowyer, K. Chang, and P. J. Flynn. An evaluation of multi-modal 2d+3d face biometrics. *IEEE PAMI*, 27(4):619-624, April 2005.

[4] K. Chang, K. W. Bowyer, and P. J. Flynn. Face recognition using 2d and 3d facial data. In *ACM Workshop on Multimodal User Authentication*, pages 25-32, December 2003.

[5] K. Chang, K. W. Bowyer, and P. J. Flynn. ARMS: adaptive rigid multi-region selection for handling expression variation in 3d face recognition. In *IEEE Workshop on FRGC Experiments*, June 2005.

[6] M. Koudelka, M. Koch, and T. Russ. A prescreener for 3d face recognition using radial symmetry and the hausdorff fraction. In *IEEE Workshop on FRGC Experiments*, June 2005.

[7] X. Lu, D. Colbry, and A. K. Jain. Three-dimensional model based face recognition. In *ICPR*, volume 1, pages 362-366, Cambridge, UK, August 2004.

[8] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview of the face recognition grand challenge. In *CVPR*, volume 1, pages 947-954, San Diego, CA, June 2005.

[9] P. J. Phillips, P. Grother, R. J. Micheals, D. M. Blackburn, E. Tabassi, and M. Bone. Face recognition vendor test 2002: evaluation report. http://www.frvt.org, March 2003.

[10] T. D. Russ, M. W. Koch, and C. Q. Little. A 2d range hausdorff approach for 3d face recognition. In *IEEE Workshop on FRGC Experiments*, San Diego, CA, June 2005.

[11] X. Yuan, J. Lu, and T. Yahagi. A method of 3d face recognition based on principal component analysis algorithm. *IEEE Symposium on Circuits and Systems*, volume 4, pages 3211-3214, May 2005.

[12] W. Zhao, R. Chellappa, and A. Rosenfeld. Face recognition: a literature survey. *ACM Computing Surveys*, 35:399-458, December 2003.

[13] R. Beveridge, D. Bolme, M. Teixeira, and B. Draper. The CSU face identification evaluation system user's guide version 5.0. http://www.cs.colostate.edu/evalfacerec, May 2003.

[14] P. Belhumeur, J. Hespanha, and D. Kriegman. Eigenfaces vs. fisherfaces: recognition using class specific linear projection. *IEEE PAMI*, 19(7):711-720, July 1997.

[15] G. Strang, *Linear Algebra and its Applications*, Thomas Learning, 1988.

[16] T. Vetter and T. Poggio. Linear object classes and image synthesis from a single example. *IEEE PAMI*, 19(7):733-742, July 1997.

[17] V. Blanz, C. Basso, T. Poggio, and T. Vetter. Reanimating faces in images and video. In *EUROGRAPHICS 2003*, volume 22, pages 641-650, Granada, Spain, September 2003.

[18] T. Heseltine, N. Pears, and J. Austin. Three-dimensional face recognition: an eigensurface approach. In *ICIP*, pages 1421-1424, October 2004.

[19] G. Passalis, I. A. Kakadiaris, T. Theoharis, G. Toderici, and N. Murtuza. Evaluation of 3d face recognition in the presence of facial expressions: an annotated deformable model approach. In *IEEE Workshop on FRGC Experiments*, San Diego, CA, June 2005.

[20] C. Hesher, A. Srivastava, and G. Erlebacher. A novel technique for face recognition using range imaging. In *Proc. Int'l Symposium on Signal Processing and Its Applications,* volume 2, pages 201-204, Tallahassee, FL, July 2003.