3D Face Recognition with the Average-Half-Face

Josh Harguess, Shalini Gupta, J. K. Aggarwal *
Computer & Vision Research Center / Department of ECE
The University of Texas at Austin
{*harguess, shalinig*}*@ece.utexas.edu, aggarwaljk@mail.utexas.edu*

Abstract

We present a promising analysis on using the pattern of symmetry in the face to increase the accuracy of three-dimensional face recognition. We introduce the concept of the 'average-half-face', motivated by the Symmetry Preserving Singular Value Decomposition. We compare face recognition results using the eigenfaces face recognition algorithm with average-half-face data and full face data in several experiments on a 3D face data set of 1126 images. We show that the results from the eigenfaces face recognition system using the average-half-face is more accurate than using the full face, only the left or right half of the face or a random choice of half of the face.

1. Introduction

We propose a technique that takes advantage of the inherent symmetry of the human face for three-dimensional (3D) face recognition. While other authors have used the inherent bilateral symmetry in face data [7, 3] to help with extracting facial profiles for recognition, none have applied this concept to 3D face recognition using subspace projection techniques, such as eigenfaces, to demonstrate the superiority of using the average-half-face instead of the full face for recognition.

The Symmetry Preserving Singular Value Decomposition (SPSVD) [5] is designed to take advantage of symmetry that is inherent to the data. The result of the SPSVD when applied to data is a symmetric approximation of the original data set. If the data is perfectly symmetric, then the SPSVD simply returns the same result as found using the Singular Value Decomposition (SVD) [1], which is simply a rank k approximation to the original data. In the case of 3D (and two-dimensional (2D)) face data comprised of images (range or intensity, respectively) of human faces in frontal poses, when the face is centered about the vertical axis of symmetry, the data is nearly symmetric. The result of the SPSVD when applied to a face image is a symmetric approximation of the original face. We store exactly one half of this data and call this the 'average-half-face'. We borrow the naming convention as introduced by Ramanathan *et al.* [4], where the authors discuss the use of the 'Half-face', which is exactly one half of the face image useful in applications of uneven illumination. We show that by using the average-half-face in an eigenfaces recognition system on 3D range images of faces, the overall accuracy of the system is significantly better than using the original full face images.

We utilize the eigenfaces face recognition method [6], which is based on the subspace projection technique known as principal components analysis (PCA), as the benchmark algorithm for our experiments by applying the technique to 3D face images. As inputs to the eigenfaces algorithm, we considered the full face, the left half of the face, the right half of the face, a random choice of either half of the face and the average-half-face.

In the paper, we first give a short description of the average-half-face. We then describe the algorithm used in the experiments. We conclude the paper by presenting results from several experiments that clearly demonstrate the accuracy gains in 3D face recognition by using the average-half-face.

2. The Average-half-face

Our method for creating the average-half-face is derived from the use of the SPSVD. For a frontal face image, that is vertically oriented, creating the average-half-face can be decomposed into two steps: centering the face within the image and then averaging the two halves of the face.

First, the face in the image is optimally centered so that the mean-squared error of the difference between

*The research was supported in part by Texas Higher Education Coordinating Board award # 003658-0140-2007.

978-1-4244-2175-6/08/$25.00 ©2008 IEEE

the two sides of the face (with the columns of one of the sides of the face reversed) is minimum. For example, in a 2D face image with equal illumination on both sides of the face, this can be achieved with a simple search for the best center of the face. For 3D range images of faces, this can also be achieved with a simple search for the best center of the face, along with the knowledge that the tip of the nose is usually the closest point (to the sensor) and is therefore easily identifiable as a starting point for the search. It is important to note that this step of centering the face image is vital for full faces, average-half-faces and any other data that use the majority of the face for face recognition, and that, for 3D faces, correcting the orientation and centering the face image is simple. Further, this processing step is done off-line and does not adversely affect the computation time of the eigenfaces method with average-half-faces as compared to full faces.

Second, the two halves of the face (the right and left half-faces) are averaged together. Note that the columns of the left half-face must be reversed so that the two half-faces are aligned before averaging.

As an example, Figure 1 (a) displays a 3D face image from our database and Figure 1 (b) displays its corresponding average-half-face. Also in Figures 1 (c) and 1 (d), we display the left and right half-faces of the same 3D face image.

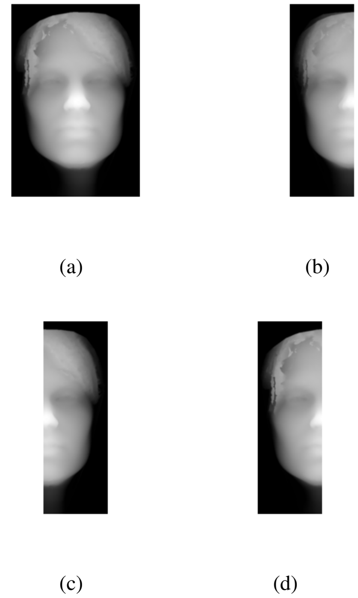

Figure 1. (a) 3D full face image; (b) its average-half-face; (c) its left half-face; and (d) its right half-face.

3. Algorithm Outline

1. Preprocess each training and test image (gallery and probe) according to the data desired (full face, average-half-face, left half-face, or right half-face). For example, to obtain average-half-faces, we split the centered face image into two (left and right) half-faces and then average the two half-faces, which is equivalent to applying a full rank SPSVD to the centered face images.

2. Perform the well-known eigenfaces algorithm [6] on the training and test sets to obtain k weights per gallery and probe image.

3. Classify the projected probe images using nearest-neighbor classification by finding the smallest Euclidean distance between the projected probe image weights and all of the projected gallery image weights.

4. Experiments

We applied the above algorithm to a database of 3D face images with frontal views of the face which was acquired at the former company Advanced Digital Imaging Research, LLC, Friendswood, TX. These range images had a resolution of 0.32 mm along the z dimension (pixel depth value) and 0.96 mm along the x and y dimensions. They were acquired using an MU-2 stereo imaging system manufactured by 3Q Technologies Ltd. (Atlanta, GA). The faces were preprocessed to remove noise as explained in [2] and centered.

In accordance with the standard practice for evaluating face recognition systems, we divided the 3D data set into three disjoint sets; train, gallery and probe sets. For the training set, we used 360 images total of 12 different subjects with (smiling and neutral) facial expressions. For the gallery set, we used one image with a neutral expression for 104 different subjects. The remaining 662 images of the 104 subjects were used as the probe set. In the probe set, the number of range images in the probe set per subject varied from 1 to 55. We used the minimum Euclidean distance in the projected subspace to classify the projected probe images to the nearest projected gallery image.

We first performed PCA on the entire 3D face data (1126 images) of full face images and on their corresponding average-half-face images to discover the amount of cumulative variance captured for each eigenvector. We did this on the full face and on the average-

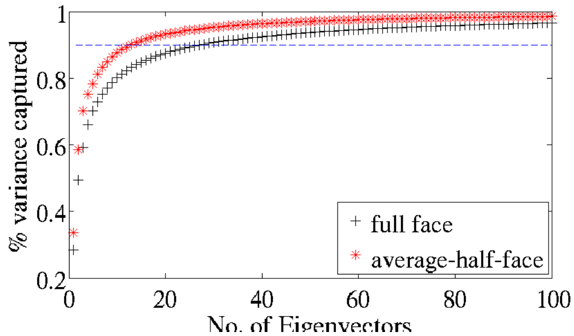

Figure 2. Cumulative variance captured of the entire 3D face database.

Table 1. Variance Captured per No. of Eigenvectors

% Variance Captured	Full Face	Average-half-face
70	5	3
75	7	4
80	10	6
85	16	8
90	28	13
95	66	29

half-face as shown in Figure 2. Table 1 lists a few examples of the number of eigenvectors needed for both the full face and the average-half-face in order to capture the same amount of cumulative variance in the data. In order to capture the same amount of variance in the data, fewer eigenvectors are required for the average-half-face than for full faces.

The next set of experiments performed involved different types of preprocessing applied to the original data. We preprocessed each image to produce an average-half-face, a left half-face (right half of the image, which contains the left half-face) and a right half-face for our experiments. We then compared the accuracy of the eigenfaces algorithm on each of these preprocessed data sets and compared them to the accuracy of the eigenfaces algorithm on the original full face data. For each of the experiments, the reported accuracy is the rank 2 accuracy, meaning that the correct classification is recorded when the correct match in the gallery is one of the two nearest images in the subspace.

Figure 3 shows the accuracy of eigenfaces using the average-half-face data compared with using the full face data. Clearly average-half-face data produces a higher recognition accuracy.

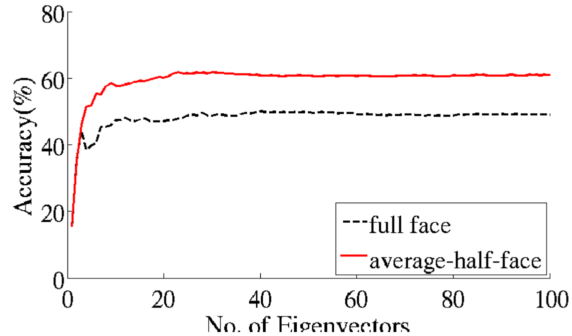

Figure 3. Rank 2 accuracy using the average-half-face.

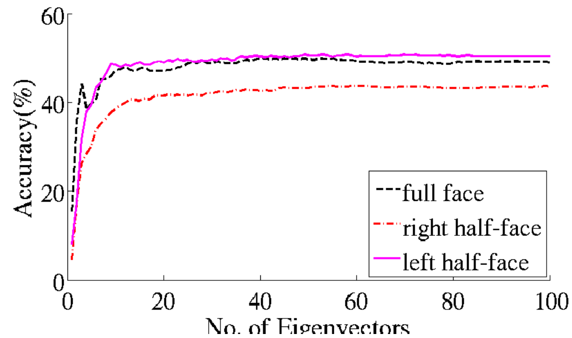

Figure 4. Rank 2 accuracy using the right and left half-face only.

For further understanding of this observed difference in the performance of the full face and the average-half-face, Figure 4 displays the accuracy of eigenfaces using only the left half-face data and only the right half-face data, respectively, as compared to using the full face data. For further examination, Figure 5 displays the results of an experiment where, in a uniformly random fashion, we chose either the right or left half-face for each image in the gallery and probe data sets.

Finally, we combine the results given by the left and right half-face data sets by using a logical 'or' rule to combine their scores. In other words, for each of the probe images, if the correct match was found in either the left or right half-face results, then we consider that a match for the combined results. The outcome of the combined results are shown in Figure 6.

5. Discussion

In our 3D face recognition experiments, the accuracy of the eigenfaces recognition when using only one

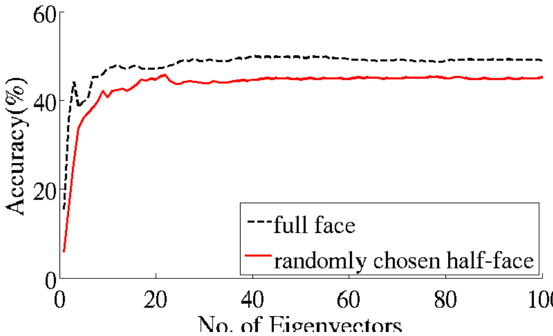

Figure 5. Rank 2 accuracy of randomly choosing the left or right half-face.

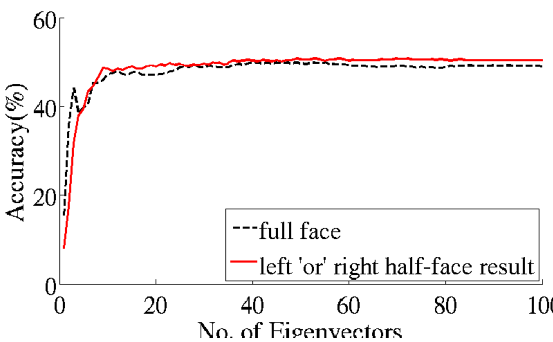

Figure 6. Rank 2 accuracy using both the left and right half-face.

half of the face (either the right or left half-face) is less than, or equivalent to, using the full face. This leads us to the conclusion that we are missing valuable information to assist in face recognition when we only consider one half of the face. It is interesting, however, that using the left half-face alone produces comparable accuracy to using the full face. Randomly choosing half of the face also performs the same or worse than using the full face. However, it is clear that in our experiments the average-half-face produces higher accuracy than the original full face data. This gain in accuracy is clearly not coming from simply using one half of the face versus the other, or even using both halves of the face simultaneously. The conclusion we draw is that the gain in accuracy has origins in the average of the two sides of the face, which is a very interesting finding. At this point in the investigation, it is not clear what other factors are behind this increase in accuracy. Therefore, further investigations are warranted into the origins of this gain in accuracy as well as the application of the average-half-face technique in two areas; other 3D subspace projection face recognition algorithms and other types of symmetric data.

The objective of the paper is to highlight the performance improvement when the inherent symmetry of the human face is exploited by using the average-half-face instead of the full face in the eigenfaces recognition algorithm. Therefore, we do not desire to compare the absolute accuracy of our system to other subspace projection algorithms. Instead, we highlight the relative accuracy of the eigenfaces algorithm between the use of the average-half-face, the full face and the left and right half-faces. This is an initial study to validate the use of the average-half-face along with a subspace projection face recognition algorithm.

6. Conclusion

We have shown that, for 3D face recognition, using the average-half-face produces better accuracy than using the original full faces when the eigenfaces method is employed. Therefore, fewer eigenvectors are required to represent each face for a given variance captured which leads to more effecient face recognition systems. It is hoped that this gain will carry over to other methods for face recognition.

References

[1] D. Bau III, L. N. Trefethen, Numerical linear algebra, Philadelphia: Society for Industrial and Applied Mathematics, pp. 25-36, 1997.

[2] S. Gupta, J. K. Aggarwal, M. K. Markey, A. C. Bovik, 3D Face Recognition founded on the Structural Diversity of Human Faces, IEEE Conference on Computer Vision and Pattern Recognition, 2007, Minneapolis, MN.

[3] G. Pan, Z. Wu, 3D Face Recognition From Range Data, International Journal of Image and Graphics, Vol.5, No.3, 2005, pp.573-593.

[4] N. Ramanathan, A. K. R. Chowdhury, R. Chellappa, Facial Similarity across Age, Disguise, Illumination and Pose, IEEE International Conference in Image Processing (ICIP), vol. 3, pp: 1999-2002, October 2004, Singapore.

[5] M. Shah, D. C. Sorensen, A Symmetry Preserving Singular Value Decomposition, SIAM Journal on Matrix Analysis and Applications, #28, 2006, pp. 749 - 769.

[6] M. Turk, A. Pentland, Eigenfaces for Recognition, Journal of Cognitive Neuroscience, Vol. 3, No. 1, 1991, pp. 71-86.

[7] L. Zhang, A. Razdan, G. Farin, J. Femiani, M. Bae, C. Lockwood, 3D face authentication and recognition based on bilateral symmetry analysis, The Visual Computer, Vol. 22, No. 1, January 2006, pp. 43-55.