Combining Face and Iris Biometrics for Identity Verification

Yunhong Wang[1], Tieniu Tan[1], Anil K. Jain[2]

[1] Center for Biometrics Authentication and Testing, National Laboratory of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences, P.O. Box 2728, Beijing 100080, P. R. China
{wangyh,tnt}@nlpr.ia.ac.cn

[2]Department of Computer Science & Engineering, Michigan State University
East Lansing, MI 48824
jain@cse.msu.edu

Abstract. Face and iris identification have been employed in various biometric applications. Besides improving the verification performance, the fusion of both of the biometrics has several other advantages such as enlarging user population coverage and reducing enrollment failure. In this paper, we make a first attempt to combine face and iris biometrics. We use two different strategies for fusing iris and face classifiers. The first strategy is to compute either an unweighted or weighted sum of the two matching distances and compare the distances to a threshold. The second strategy is to treat the matching distances of face and iris classifiers as a two-dimensional feature vector and use a classifier such as the Fisher's discriminant analysis or a neural network with radial basis function (RBFNN) to classify the vector as being genuine or an impostor. We compare the results of the combined classifier with the results of the individual face and iris classifiers. Experimental results show the validity of the proposed fusion strategies.

Keywords. fusion, face, iris, verification

1 Introduction

With increasing concern and demands on security, automatic identity verification systems based on biometrics have become more widespread and more efficient. Several airports are now equipped with biometric products. Face and iris recognition systems are among the top choices: face recognition is the most friendly and non-invasive way of recognizing a person whereas iris recognition is one of the most accurate biometrics [1][2]. However, there are a number of practical issues that still need to be solved with both systems. The accuracy of face recognition is extremely sensitive to illumination, pose and expression [3]. In many applications, face identification systems must be robust to these variations. In the case of iris recognition, the user must be cooperative. Further, iris images must meet stringent quality criteria, so the images of poor quality (e.g., iris with large pupil, or off-center images) are automatically rejected at the time of acquisition. Consequently, several attempts may be necessary to acquire the iris image, which not only delays enrollment and verification, but is also an annoyance to the user. The rate of rejection of poor quality images is termed "enrollment failure rate". Like any other biometric, the iris can change (e.g., as a result of eye disease), in which case, even a very accurate iris based identification system can fail.

Some of the above problems can only be solved, or at least their impact can be reduced, by fusing several biometric identification systems, such as face and iris recognizers. It is well known that by fusing several classifiers, the overall error rate (the false accept rate and the false reject rate) can go down [4]. This fusion also reduces spoof attacks on the biometric system. The population coverage of a combined classifier is, in general, larger than the coverage of a component classifier, regardless of the accuracy of the latter; people with various disabilities may only be able to provide certain biometrics and not others. Thus, combining classifiers increases the number of people that can be identified. While it is true that a combined classifier requires that the user provide several biometrics during the acquisition stage, the combination of face and iris allows for simultaneous acquisition of these two images (for example, by using two cameras, one for face and the other for iris). Thus, in this particular case, no additional inconvenience is introduced. Finally, the use of face recognizer in addition to the iris classifier, may allow people with imperfect iris images to enroll, reducing the enrollment failure rate.

A substantial amount of work has been done on the combination of multiple classifiers [5][6][7][8][9][18][19]. Most of these studies focus on fusing 'weak' classifiers, for the purpose of increasing the overall performance. However, the advantages outlined above warrant the combination of available 'strong' classifiers, for purposes other than increased performance. It has been shown that the simple sum rule gives very good accuracy in combining multiple biometric systems [5][6][7][19]. But this holds true only when the component classifiers are of similar accuracy. Can the sum rule still give very good accuracy when it is used to combine a 'weak' classifier with a 'strong' classifier? Not much work has been done in this direction. In this paper we attemp to answer this question by developing a combined face-iris verification system in which the component face and iris classifiers are of very different accuracy. The combined system overcomes a number of inherent difficulties of the component classifiers. The strategies for developing fusion rule are based on learning and non-learning methods. The accuracy of the fused systems is compared with that of the individual face and iris classifiers. Moreover, accuracies resulting from different strategies are compared and analyzed. A preliminary version of this work appears in [29].

The remainder of the paper is organized as follows. Section 2 describes the operation of face and iris recognition systems. Section 3 describes the fusion method. Section 4 presents experimental results when using the proposed approach. A detailed analysis of experimental results and conclusions are given in Section 5.

2 Face Verification and Iris Verification

2.1 Face verification

Among various face recognition algorithms, appearance-based approaches are the most popular. Some of the representative methods are PCA [10], ICA [11] and LDA [12]. Here, we use one of the most popular algorithms, the Eigenface method as the face matcher. Since our focus is on the fusion method, some of the details of the face recognition system are not described. As we know, a face recognition system includes face detection, preprocessing, feature extraction and matching [3]. We assume that all the faces are localized and aligned before they are recognized.

Our face verification system is a standard principal component analysis classifier (PCA) [10]. Let the ith training face image be represented as an N-dimensional vector X_i, i=1,2,…n. The scatter matrix S of all the training samples can be computed as

$$S = \sum_i (X_i - m)(X_i - m)^T \ , \tag{1}$$

where m is the mean vector.

The principal directions of S are the eigenvectors corresponding to the M largest eigenvalues of S, M<<N. For each input face image X, we obtain a feature vector Y by projecting X onto the subspace generated by the principal directions, according to the following equation:

$$Y = W^T X \ . \tag{2}$$

The projection matrix W is the matrix whose rows are the principal directions. The dimension of the feature vector Y is equal to the number of "significant" principal directions, which is, in general, a few orders of magnitude smaller than the size of the original input vectors. Images are then compared by means of their corresponding feature vectors.

2.2 Iris Verification

Iris recognition has received increasing attention due to its high reliability for personal identification [12]. The human iris, an annular part between the black pupil and the white sclera, has an extraordinarily rich structure and provides many interlacing minute characteristics such as freckles, coronas, stripes, furrows, crypts and so on. These visible characteristics, generally called the texture of the iris, are unique to each subject [13][14][15][16]. Individual differences that exist in the development of anatomical structures in the body result in such uniqueness, which leads to high reliability for personal identification. There are four main approaches to iris representation: phase information coding [13][16], zero-crossings representation [24], shape description [25] and texture analysis [14-15][25-28]. The iris recognition system employed in this paper is based on an efficient algorithm that characterizes the critical points of local variations.

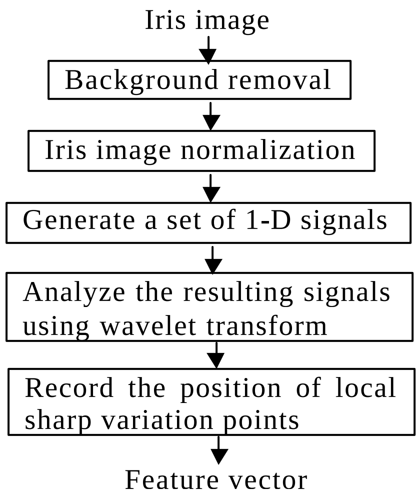

Figure 1. Major steps in iris feature extraction.

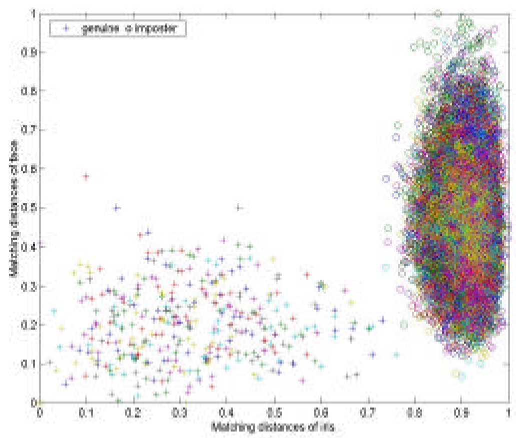

Figure. 2. Two-dimensional distribution of matching distances based on face and iris.

Figure 1 illustrates the main steps of our method [17]. First, the background in the iris image is removed by localizing the iris. Then, the annular iris region is normalized to a rectangular block of a fixed size. After lighting correction and image enhancement, a set of 1-D intensity signals containing the main spatial variations of the original iris for the subsequent feature extraction is constructed. We record the position of local sharp variation points in each intensity signal as features by wavelet analysis. A matching scheme based on Exclusive OR calculation is proposed. Further details on the iris recognition system may be found in [17].

3 Fusion of Classifiers

The verification problem using a combination of classifiers can be formulated as follows. Classify a test sample S into one of the following two classes: w_0 (genuine) or w_1 (impostor). If x_1 and x_2 are the outputs of the component classifiers, then

$$\text{Assign} \quad S \to w_j \quad \text{if} \quad P(w_j|x_1, x_2) = \max_{k=0}^{1} P(w_k|x_1, x_2) \quad j=0,1, \tag{3}$$

where $P(w_k|x_1, x_2)$ denotes the posteriori probability of w_k given x_1 and x_2.

There are two different strategies that we employ for fusing the classifiers. One strategy is to compute either an unweighted or weighted sum of x_1 and x_2, and to compare the result to a threshold. The second strategy is to treat x_1 and x_2 as a two-dimensional feature vector. We then use a classifier such as Fisher's discriminant analysis and a neural network with radial basis function (RBFNN) to decide whether the vector (x_1, x_2) represents a genuine sample or an impostor. We present the details of these methods below.

3.1 Sum rule

The sum rule [5][6][7][8] has been extensively investigated and it is the most straightforward fusion strategy. It operates directly on the soft outputs (matching distances) of individual classifiers for each class hypothesis. We first use equal weights in the fusion rule. The matching distances of individual classifiers are normalized linearly in advance.

According to Kittler's fusion framework, the sum rule can be expressed as [5]

$$\text{Assign} \quad S \to w_j \quad \text{if} \quad \sum_{i=1}^{2} P(w_j|x_1, x_2) = \max_{k=0}^{1} \sum_{i=1}^{2} P(w_k|x_1, x_2) \quad j=0,1, \tag{4}$$

where i is the individual classifier.

3.2 Weighted sum rule

The performances of different classifiers are different, so it is necessary to use different weights to combine the component classifiers. Here, we utilize the EER (Equal Error Rate, where FAR=FRR) of each classifier as weights.

The weighted sum rule is defined as:

Assign S $\rightarrow w_j$ if $\sum_{i=1}^{2} W_i P(w_j | x_1, x_2) = \max_{k=0}^{1} \sum_{i=1}^{2} W_i P(w_k | x_1, x_2)$ j=0,1 (5)

where i is the individual classifier, W_i is the weight assign to classifier i.

The weights are calculated as follows:

$$W_i = \frac{1-2E_i}{2-(2E_i+2E_t)} \qquad \text{i=1,2, t=1,2 ,} \quad i \neq t \tag{6}$$

where i and t represent the face and iris verification system, respectively, E_i is the EER of classifier i. Obviously, $\sum_{i=1}^{2} W_i = 1$.

In fact, the accuracy of the verification result depends on two things: the accuracy of the individual biometric system and the typicality of the subject that means whether the feature extraction method is appropriate to the specific user. It is true that a good recognition system can succeed for most of the subjects. But for a specific user, the accuracy of verification is also dependent on the typicality. So it is necessary to assign different weights to each user to obtain high verification accuracy. Jain and Ross [9] investigated this problem and found the user-specific weights by 'exhaustive' search. This approach is not feasible when the number of users is large. Here we set the weights according to the performance of the verification results of the subject with different classifiers. That is,

$$W_{ui} = \frac{1-2E_{ui}}{2-(2E_{ui}+2E_{ut})} \quad , \tag{7}$$

where W_{ui} is the weight assigned to classifier i for user u, E_{ui} is the equal error rate of user u with classifier i. Obviously, for the two-classifier case considered here, $\sum_{i=1}^{2} W_{ui} = 1$.

3.3 Fisher discriminant analysis

If we treat the face and iris matcher outputs x_1 and x_2 as a feature vector X=(x_1, x_2), then we can use any of the known classifiers to determine the separation boundary between the impostor and genuine samples. If we seek a linear boundary, then it is known that the line that, under the Gaussian assumption with equal covariance matrices, best separates the two classes can be computed as

$$W^* = S^{-1}(m_1 - m_2), \tag{8}$$

where m_1, m_2 and S are defined as follows.

$$S = \sum_{x \in D_0} (x-m_1)(x-m_1)^T + \sum_{x \in D_1} (x-m_2)(x-m_2)^T \tag{9}$$

$$m_1 = \frac{1}{n} \sum_{x \in D_0} x_1 \qquad m_2 = \frac{1}{n} \sum_{x \in D_1} x_2 \,.$$

D_0 is the space of genuine samples and D_1 is the space of impostor samples.

The Fisher based fusion method can be expressed by the following equations:

$$x \in w_1 \text{ if } Y \geq Y_0, \quad \text{else } x \in w_0 \text{ , where } Y = W^{*T} X \;. \tag{10}$$

Figure 2 shows the distribution of the genuine and impostor matching distances for face and iris. Clearly, some users cannot be identified based solely on the matching distance of face or iris. The experimental results in Section 4 show that the verification rates are improved by fusion with Fisher rule.

3.4 RBF Neural Network based fusion method

Again, forming a vector (x_1, x_2) from the individual outputs of the face and iris classifiers, we use a neural network with radial basis function (RBFNN) for classification. We chose RBFNN over other types of multi-layer perceptron neural networks, because it had the best performance in our experiments. The other reason is that this neural network can learn either the positive samples or the negative samples. This property makes RBFNN suitable for verification. We use 2 nodes in the input layer, 10 nodes in the hidden layer and 1 node in the output layer.

The output of the jth hidden node in a RBF neural network can be expressed as [20]:

$$\mathbf{y}_h = \Phi(\|\mathbf{X}_k - \mathbf{C}_j\|), \text{ j=1,2,... } N_0 , \tag{11}$$

where $\mathbf{X}_k$ is a 2-dimensional input vector, $\mathbf{C}_j$ is the jth RBF neural network's center, N_0 is the number of hidden units, and $\Phi(\cdot)$ is a nonlinear, radial symmetric function whose center is C_j.

The output of the ith output unit of the RBFN is:

$$z_{ki} = \sum_h w_{ih} \Phi(\|\mathbf{X}_k - \mathbf{C}_\mathbf{j}\| + w_{k0} \,. \tag{12}$$

For identity verification problem, the number of genuine samples is far fewer than the number of impostors (as for a given subject, all other individuals can be treated as imposters). As the typicality of each subject is different, matching distances of genuine and impostor of each subject are different. It is necessary to design different RBF classifiers to fuse the individual matcher results. Notice that since we are dealing with a verification problem, we can build an individual neural network for each subject in the database. This is because, for verification, an unknown individual must claim his identity first, so we would know which neural network to use.

4. Experiments Results

4.1 Databases

We have used two databases in our experiments, each containing face images and iris images. Aside from the number of subjects in each database, the main difference is that the quality of the iris images (the iris images in Database 1 are of very high quality, whereas those in Database 2 are of somewhat lower quality).

4.1.1 Database 1

We collected the face images in Database 1 from the ORL [21], MIT [22], Yale [23] and NLPR databases. While the first three are well known public domain face databases, the NLPR face database consists of face images taken in our lab at two different time instants (NLPR is for National Laboratory of Pattern Recognition). Examples of typical face images from the NLPR database are shown in Figure 3. The ORL database contains 40 subjects and 400 images. The MIT database contains 16 subjects and 432 images. The Yale database contains 15 subjects and 165 images. The NLPR database contains 19 subjects and 266 images. For each subject we selected 5 images, yielding a total of 450 face images for 90 subjects. The integrated face database is composed of faces with reasonable variations in expression, pose and lighting.

There are no public domain iris image databases. The iris images used in our system come from the *NLPR iris database*. The database includes 2,096 iris images corresponding to 210 subjects, captured by an iris acquisition system developed at the NLPR [15]. There are at least 5 images for each eye. Each eye of a person represents a different class since the iris images of the left and right eyes are known to be different. Since not every individual provided iris images of both eyes, there are 303 different classes from 210 subjects. The images were acquired during two different sessions, one month apart. Figure 4 shows some sample iris images from the NLPR database.

Obviously, the ORL, MIT and Yale face databases do not come with corresponding iris images, so to each face image, we assign an arbitrary (but fixed) iris class. In this way, we obtain a database of 90 subjects, with 5 face images and 5 iris images per subject.

The protocal of Database 1 is that for face verification, we use the leave-one-out method to establish the Eigenspace. This means that 4 faces of each subject can be used for training and one image for testing. We replaced the test image with one of the training images, and repeated this procedure several times. The number of matches for constructing the genuine distribution is 5*90 and the number of matches for constructing the impostor is 5*90*89. For iris verification we use 4 irises of each subject for training and one iris image for testing, similar to procedure used in face verification. As all of the matching distances obtained in this way can be regarded as results of testing data, the results are divided into two parts for the learning based methods employed for fusion. For each subject, we utilized 3 samples (here are the matching distances of face and iris) for training and 2 for testing.

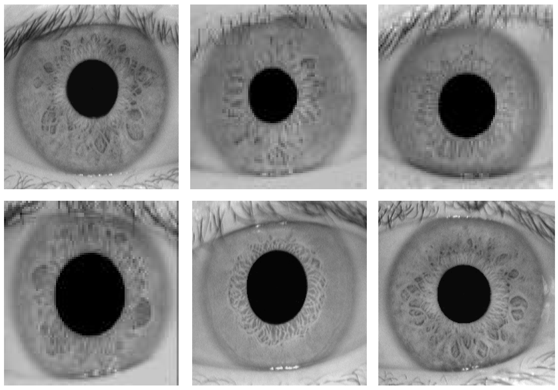

Figure.3. Sample face images in the NLPR database. Figure. 4. Sample iris images in the NLPR database.

4.1.2 Database 2

To illustrate the enrollment failure and its effect on the overall verification accuracy, we use some of the iris images with poor quality that would normally be rejected in an operational iris verification system. Samples of poor quality iris images are shown in Figure 5. We use 40 subjects and 400 iris images, 10 images per subject. The face database is the ORL database that includes 40 subjects and 400 images with 10 images for each subject.

Figure. 5. Samples of poor quality iris images in Database 2 due to eyelid/eyelash occlusion

The protocal of Database 2 is that, for each subject, we use 5 images for training and 5 images for testing in the individual classifiers. Then, the training and testing patterns are exchanged. The training phase is repeated. This means that for face verification, we use 200 samples to establish the eigenspace as the training data. For iris verification we use 5 images as training images of each person, 5 for testing, and then the training and testing sets are exchanged. As all of the matching distances obtained in this way can be regarded as results of testing data, the results are divided into two parts for the learning based methods employed for fusion. We utilized 6 samples for training and 4 samples for testing per person.

4.2 Experimental results

4.2.1 Results on Database 1.

The verification rate of the stand-alone iris classifier is very high. Most of the FAR and FRR rates at various thresholds are zero. It is difficult to show and compare the FAR and FRR rates using the ROC curve. Therefore, we simply calculated the total error rate (i.e., FAR+FRR). The results based on the total error rate of all of the verification systems (include combined systems) are calculated. The result show that the RBF classifier get the fewest total error rate.

The distribution of matching distances of iris and face are shown in Figures 6 (a) and (b), respectively. The distributions of the matching distances for the fused systems are shown in Figures 6 (c), (d) and (e). Since the outputs of combined system fusing with RBFNN are near 0 or 1, the distribution of genuine and impostor can not be shown in the Figure 6. We just compared its total error to others. It is apparent from the distribution of matching distances that fusion improves the separation of genuine scores from the impostor scores. As the ROC curves cannot be compared based on the reasons mentioned above, we analyze the range of thresholds that result in very high accuracy. The range of threshold values that can be selected for high verification rates is larger for the fused matcher compared to the individual matchers.

There are a number of thresholds that can result in zero error rates for iris and combined systems. The advantage of fusion is that we can get a larger range of operating points (thresholds) with high accuracy. In an operational system, the thresholds are set using the training data. Sometimes, the training data is not representative of the underlying population and the impostor data cannot be acquired to learn the threshold.

A larger range of thresholds that result in high accuracy means that the system is more robust.

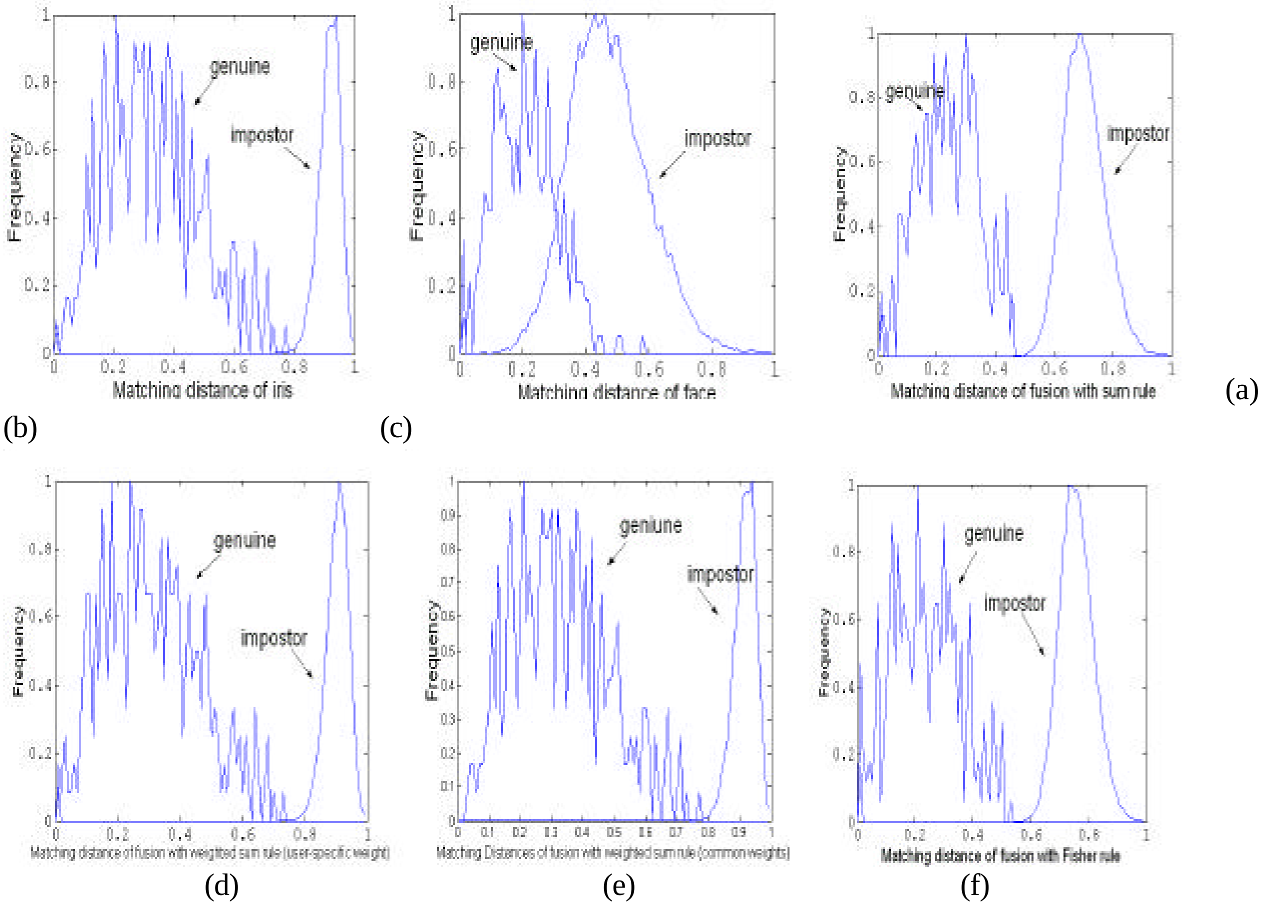

Figure. 6. Distribution of matching distances of (a) iris, (b) face, (c) sum rule, (d) weighted sum rule (user specific weights) (e) weighted sum rule (common weights), (f) Fisher rule.

4.2.2 Results on Database 2

The performance of the iris verification and the combined system on Database 2 is not as high as on Database 1 because of the poor quality iris images in Database 2. Now we can compare the matchers by means of the ROC curves shown in Figure 7.

For comparison, we give the minimum error rates of the various verification systems in Table 1. We can see that the RBF fusion method achieves the highest verification accuracy.

Although the iris system has high accuracy, it is not perfect. Figure 8 and Figure 9 show some samples on which the stand-alone iris classifier fails. The difference between the template and test sample in Figures 8 and 9 is large.

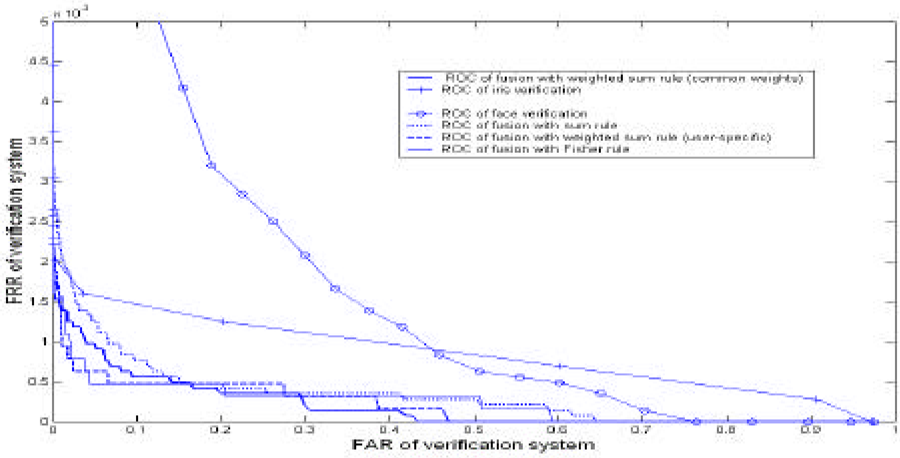

Figure. 7 . ROC curves of verification systems (at a fine resolution).

Table 1 Minimum error rates of verification systems

Classification System	Minimum of error rates	Classification System	Minimum of error rates
Iris	0.0029	Fusion with weighted sum rule	0.0031(common)
			0.0028(user-specific)
Face	0.0169	Fusion with Fisher rule	0.0025
Fusion with sum rule	0.0037	Fusion with RBF	0.0024

Fig. 8 (a)

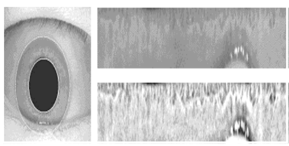

Fig. 9 (a)

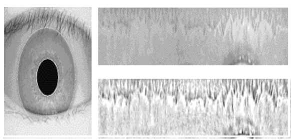

Fig 8. (b)

Fig. 9 (b)

Fig. 10

Figure 8. Mismatch caused by eyelids and eyelashes covering the iris. (a) templates (b) test sample

Figure 9. Mismatch caused by different pupil size in the template. (a) templates (b) test sample.

Figure 10. Mismatched faces

Figure 8 and Figure 9 show some of the samples on which the standalone iris classifier failed, but they were correctly recognized by the combined system. Since the combined system does not depend on a single modality, some of the subjects that failed in the individual verification system can be verified correctly by the combined system. An operational iris verification system would reject poor quality images like the ones in Figures 8 and 9 in the acquisition stage to maintain a high accuracy and require the user to try again to get a better iris image. Our system demonstrates that fusion is a way to decrease the enrollment failure rate. At the same time, a larger number of subjects that were misclassified by the stand-alone face verification system were correctly identified by the combined system. Samples of mismatched faces are shown in Figure 10.

Based on the above experimental results, we can draw the following conclusions:

1. Learning based methods are suitable to combine classifiers when their accuracy is not similar. Here we fuse an iris verification system that has a high accuracy and a face verification system that has a considerably lower accuracy. All the training (learning) based methods, including the RBF neural network, the weighted sum rule and the Fisher rule are valid for combining face and iris verification systems and perform better than methods without learning (e.g, sum rule).
2. The RBF neural network based fusion method provides the best result on the two databases because we utilize it as a user-specific method and train a specific RBF neural network for each user. The weights of neural network are set by learning the positive (genuine) and negative (impostor) samples. The weighted sum rule that uses user specific weights obtains similar accuracy as the RBF neural network. The reason for this is that the accuracy of the verification result depends on two things: (i) accuracy of the individual biometric system, and (ii) the typicality of the subject which means whereas the feature extraction method is appropriate to the specific user. User-specific based method considers the above two aspects and is more accurate.
3. The results on Database 1 show that the range of thresholds that lead to high verification rates of combined systems is larger than that of individual system. Since the thresholds be set only by training, and as the training data is not enough (in fact, real impostor data can not be acquired, even the genuine data suffer from the small sample size), the thresholds are not accurate. Enlarging the range of thresholds that result in high accuracy means that the recognition system is more robust. Thus the fusion based verification system is more robust that the individual verification systems.
4. The results on Database 2 show that the combined system can accept some iris images with poor

quality. An iris recognition system requires more user cooperation than with a face recognition system. To maintain a high accuracy of iris verification, an 'enrollment failure rate' is set to ensure that all the acquired iris images are of good quality. This is not convenient to the users. The combined system can decrease the 'enrollment failure rate' (as it can accommodate some low quality iris images) and obtain high accuracy of identity verification. Some of the fusion methods show their higher performance than the individual verification systems.

5 Conclusions

We have designed and built an identity verification system based on the fusion of face and iris data. The significance of fusing these two biometrics is more than the improvement in verification accuracy. Enlarging user population coverage and reducing enrollment failure, two very important factors in practical applications, are additional reasons for combining face and iris for verification. We have used two strategies for fusion: (i) weighted/unweighted summation of the outputs (x_1, x_2) of the standalone classifiers, and (ii) treating (x_1, x_2) as a 2D feature vector, in which case we used the Fisher discriminant and a neural network classifier. Combined systems show more robust performance than the iris verification system and the face verification system alone. Fusion based on the RBF neural network produced the highest verification accuracy. In general, the experimental results also show that the learning based methods perform better when they are used to fuse a 'strong' classifier and a 'weak' classifier. Experimental results have further demonstrated that the enrollment failure rate of stand-alone systems can be decreased by fusion, while maintaining a high accuracy.

ACKNOWLEDGMENTS

Thanks to Silviu Minut for a careful reading of this paper. Thanks to Li Ma and Yuchun Fang for providing the iris and face recognition algorithms. Thanks to Arun Ross for his suggestions regarding fusion approaches. This research was supported by a grant from the NSF IUC on Biometrics (CITeR) at West Virginia University, the Chinese National Hi-Tech R & D program Grant No. 2001AA114180), and the NSFC (No.69825105).

References

1. A.K. Jain, R.M. Bolle and S. Pankanti, Eds., *Biometrics: Personal Identification in a Networked Society*, Norwell, MA: Kluwer, 1999.

2. T. Mansfield, G. Kelly, D. Chandler and J. Kane, Biometric Product Testing Final Report, issue 1.0, National Physical Laboratory of UK, 2001.
3. W. Zhao, and R. Chellappa, A. Rosefeld, and P.J. Phillips, Face Recognition: A Literature Survey, Technical Report, Computer Vision Lab, University of Maryland, 2000.
4. L. Hong and A.K. Jain, "Integrating Faces and Fingerprints for Personal Identification", IEEE Trans. PAMI, 20(12): 1295-1307, 1998.
5. Josef Kittler, Mohamad Hatef, Robert P. W. Duin, and Jiri Mates, "On Combining Classifiers", IEEE, Trans. PAMI, 29(3): 226-239, 1998.
6. Fabio Roli, Josef Kittler, Giorgio Fumera and Daniele Muntoni, " An Experimental Comparison of Classifier Fusion Rules for Multimodal Personal Identity Verification Systems", Proc. Multiple Classifier Systems, Sringer-Verlag, LNCS 2364, 325-336, 2002.
7. Josef Kittler and F. M. Alkoot, "Relationship of Sum and Vote Fusion Strategies, Proc. Multiple Classifier Systems, Springer-Verlag, LNCS 2096, 349-358, 2001.
8. Luis Alexandre, Aurelio Campilho, and Mohamed Kamel, "Combining Independent and Unbiased Classifiers Using Weighted Average", Proc. 15th ICPR, Vol. 2, 495-498, 2000.
9. A. K. Jain and A. Ross, "Learning User-specific Parameters in a Multibiometric System", Proc. IEEE Inter. Conf. on Image Processing (ICIP), Vol. 1, 57-60, 2002.
10. M. Turk and A. Pentland, "Eigenfaces for Recognition", Journal of Cognitive Neuroscience, 3(1):71-86, 1991.
11. M. S. Bartlett, H. M. Lades and T. J. Sejnowski, "Independent Component Representations for Face Recognition", Proc. SPIE, Vol. 3299, 528-539, 1998.
12. P. N. Belhumer, J. P. Hespanha and D. J. Kriegman, "Eigenfaces vs. Fisherfaces: Recognition Using Class Specific Linear Projection", IEEE Trans. PAMI, 19(7): 711-720, 1997.
13. J.G. Daugman, " High Confidence Visual Recognition of Persons by a Test of Statistical Independence", IEEE Trans. PAMI, 15(11): 1148-1161, 1993.
14. Y. Zhu, T. Tan, and Y. Wang, " Biometric Personal Identification Based on Iris Patterns", Proc. 15th ICPR, Vol. 2, 805-808, 2000.
15. L. Ma, Y. Wang, and T. Tan, "Iris Recognition Using Circular Symmetric Filters", Proc. 16th ICPR, Vol.2, 414-417, 2002.
16. J.G. Daugman, "Statistical Richness of Visual Phase Information: Update on Recognizing Persons by Iris Patterns", International Journal of Computer Vision, 45(1): 25-38, 2001.
17. Li Ma, Tieniu Tan and Yunhong Wang, " Efficient Iris Recognition by Characterizing Key Local Variations", IEEE Trans. on Image Processing (accepted subject to minor revision).
18. M. Skurichina and R.P.W. Duin, " Bagging, Boosting and the Random Subspace Method for Linear Classifiers", Pattern Analysis and Applications, 5(1): 121-135, 2002.
19. D.M.J. Tax, M. van Breukelen, R.P.W. Duin and J. Kittler, "Combining Multiple Classifiers by Averaging or by Multiplying?", Pattern Recognition, 33(9): 1475-1485, 2000.
20. C. M. Bishop, *Neural Networks for Pattern Recognition*, Oxford University Press, UK, 1995.
21. Ferdinando Samaria and Andy Harter, "Parameterization of a Stochastic Model for Human Face Identification", Proc. 2nd IEEE Workshop on Applications of Computer Vision, Sarasota FL, 1994.
22. ftp://whitechapel.media.mit.edu/pub/images.
23. Yale University Face Database, http://cvc.yale.edu/projects/yalefaces.html.
24. J.G.Daugman, "Wavelet Demodulation Codes, Statistical Independences, and Pattern Recognition", Proc. 2nd IMA-IP, Institute of Mathematics and its Applications, 244-260, London: Horwood, 2000.
25. W.W.Boles and B. Boashah, "A Human Identification Technique Using Images of the Iris and Wavelet Transform", IEEE Trans. Signal Processing, 46(4): 1185-1188, 1998.
26. L. Ma, T. Tan, D. Zhang and Y. Wang, "Local Shape Description Based Iris Recognition", (Submitted to Pattern Recognition).
27. R.P. Wildes, J.C. Asmuth, et al., "A Machine-vision System for Iris Recognition", Machine Vision and Applications, 9: 1-8, 1996.
28. L. Ma, Y. Wang and T. Tan, "Iris Recognition Based on Multichannel Gabor Filter", Proc. Fifth Asian Conference on Computer Vision, Vol. I, 279-283, 2002.

29. Y. Wang, T. Tan and A. Jain, “Combining Face and Iris Biometrics for Identity Verification”, to appear on Proc.AVBPA’2003, 2003.