Proceedings of Biometric Authentication Workshop, LNCS 3087, pp. 259-269,
Prague, May 2004

Integrating Faces, Fingerprints, and Soft Biometric Traits for User Recognition

Anil K. Jain, Karthik Nandakumar, Xiaoguang Lu, and Unsang Park

Department of Computer Science and Engineering
{jain,nandakum,lvxiaogu,parkunsa}@cse.msu.edu
Michigan State University, MI - 48824, USA

Abstract. Soft biometric traits like gender, age, height, weight, ethnicity, and eye color cannot provide reliable user recognition because they are not distinctive and permanent. However, such ancillary information can complement the identity information provided by the primary biometric traits (face, fingerprint, hand-geometry, iris, etc.). This paper describes a hybrid biometric system that uses face and fingerprint as the primary characteristics and gender, ethnicity, and height as the soft characteristics. We have studied the effect of the soft biometric traits on the recognition performance of unimodal face and fingerprint recognition systems and a multimodal system that uses both the primary traits. Experiments conducted on a database of 263 users show that the recognition performance of the primary biometric system can be improved significantly by making use of soft biometric information. The results also indicate that such a performance improvement can be achieved only if the soft biometric traits are complementary to the primary biometric traits.

1 Introduction

Biometric systems recognize users based on their physiological and behavioral characteristics [1]. Unimodal biometric systems make use of a single biometric trait for user recognition. It is difficult to achieve very high recognition rates using unimodal systems due to problems like noisy sensor data and non-universality and/or lack of distinctiveness of the chosen biometric trait. Multimodal biometric systems address some of these problems by combining evidence obtained from multiple sources [2]. A multimodal biometric system that utilizes a number of different biometric identifiers like face, fingerprint, hand-geometry, and iris can be more robust to noise and alleviate the problem of non-universality and lack of distinctiveness. Hence, such a system can achieve a higher recognition accuracy than unimodal systems. However, a multimodal system will require a longer verification time thereby causing inconvenience to the users.

It is possible to improve the recognition performance of a biometric system without compromising on user-friendliness by utilizing ancillary information about the user like height, weight, age, gender, ethnicity, and eye color. We refer to these traits as soft biometric traits because they provide some information about the individual, but lack the distinctiveness and permanence to sufficiently differentiate any two individuals (see Figure 1 for examples of soft biometric traits). The soft biometric traits can either be continuous or discrete. Traits such as gender, eye color, and ethnicity are discrete

in nature. On the other hand, traits like height and weight are continuous variables. Heckathorn et al. [3] have shown that a combination of soft attributes like gender, race, eye color, height, and other visible marks like scars and tattoos can be used to identify an individual only with a limited accuracy. Hence, the ancillary information by itself is not sufficient to recognize a user. However, soft biometric traits can complement the traditional (primary) biometric identifiers like fingerprint and hand-geometry and hence improve the performance of the primary biometric system.

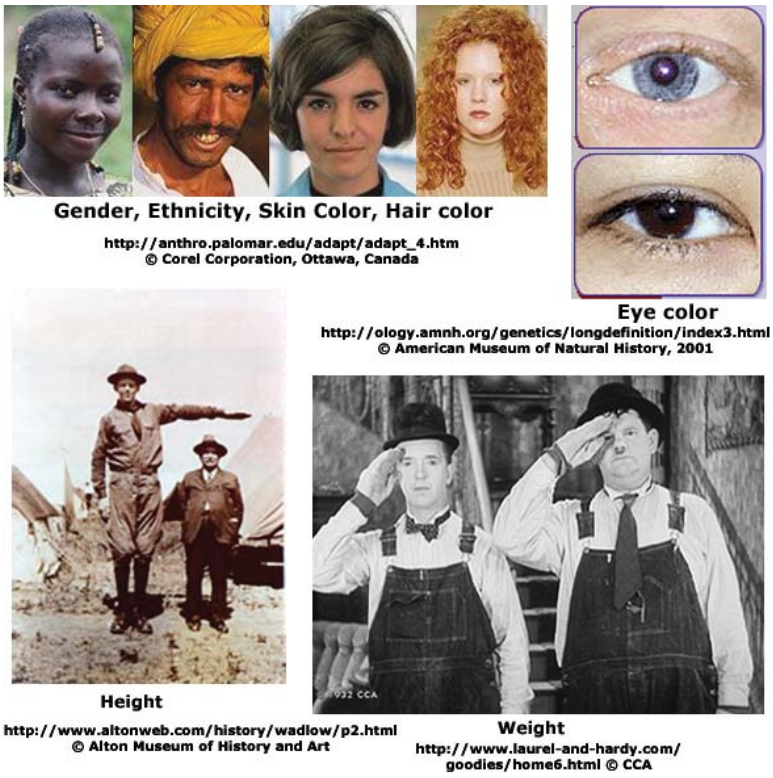

Fig. 1. Examples of soft biometric traits.

In order to utilize soft biometrics, there must be a mechanism to automatically extract these features from the user during the recognition phase. As the user interacts with the primary biometric system, the system should be able to automatically extract the soft biometric characteristics like height, weight, age, gender, and ethnicity in a non-obtrusive manner without any interaction with the user. In section 2 we present some of the methods that could be used for automatic extraction of the soft biometric information. Section 3 describes our framework for the integration of soft biometrics with the primary biometric system. The objective of this work is to analyze the impact of

introducing soft biometric variables like gender, ethnicity, and height into the decision making process of a recognition system that uses faces and fingerprints as the primary biometric traits. The experimental results presented in section 4 give an insight on the effects of different soft biometric variables on the recognition performance.

2 Automatic Extraction of Soft Biometric Characteristics

Soft biometric characteristics like gender, ethnicity, and age could be derived from the facial image of the user. Several studies have attempted to identify the gender, ethnicity, and pose of the users from their facial images. Gutta et al. [4] proposed a mixture of experts consisting of ensembles of radial basis functions for the classification of gender, ethnic origin, and pose of human faces. They also used a SVM classifier with RBF kernel for gating the inputs. Their gender classifier classified users as either male or female with an average accuracy rate of 96%, while their ethnicity classifier classified users into Caucasian, South Asian, East Asian, and African with an accuracy of 92%. These results were reported on good quality face images from the FERET database that had very little expression or pose changes. Based on the same database, Moghaddam and Yang [5] showed that the error rate for gender classification can be reduced to 3.4% by using an appearance-based gender classifier that uses non-linear support vector machines. Shakhnarovich et al. [6] developed a demographic classification scheme that extracts faces from unconstrained video sequences and classifies them based on gender and ethnicity. Their demographic classifier was a Perceptron constructed from binary rectangle features. The learning and feature selection modules used a variant of the AdaBoost algorithm. Their ethnicity classifier classified users as either Asian or non-Asian. Even under unconstrained environments, they showed that a classification accuracy of more than 75% can be achieved for both gender and ethnicity classification. For this data, the SVM classifier of Moghaddam and Yang had an error rate of 24.5% and there was also a notable bias towards males in the classification (females had an error rate of 28%). Balci and Atalay [7] reported a classification accuracy of more than 86% for a gender classifier that uses PCA for feature extraction and Multi-Layer Perceptron for classification. Jain and Lu [8] proposed a Linear Discriminant Analysis (LDA) based scheme to address the problem of ethnicity identification from facial images. The users were identified as either Asian or non-Asian by applying multiscale analysis to the input facial images. An ensemble framework based on the product rule was used for integrating the LDA analysis at different scales. This scheme had an accuracy of 96.3% on a database of 263 users (with approximately equal number of users from the two classes).

Automatic age determination is a more difficult problem due to the very limited physiological or behavioral changes in the human body as the person grows from one age group to another. There are currently no reliable biometric indicators for age determination [9]. Buchanan et al. [10] have been studying the differences in the chemical composition of child and adult fingerprints that could be used to distinguish children from adults. Kwon and Lobo [11] present an algorithm for age classification from facial images based on cranio-facial changes in feature-position ratios and skin wrinkle analysis. They attempted to classify users as “babies”, “young adults”, or “senior adults”. However, they do not provide any accuracy estimates for their classification scheme.

One can hope that age determination systems providing a reasonable estimate of the age of a person would be available in the near future.

The weight of a user can be measured by installing a weight sensor at the place where the users stand while providing the primary biometric. The height can be estimated from a sequence of real-time images obtained when the user moves into the view of the camera. Figure 2 describes a mechanism for simultaneous extraction of the height information and the facial image of a user. In this setup we assume that the position of the camera and the background scene are fixed. The background image (Figure 2(a)) is initially stored in the system. Two markers are placed in the background for calibration. The first marker is placed at a height H_{low} above the ground and the second marker is placed at a distance H_{ref} above the first marker. The vertical distance between the two markers in the background image is measured as D_{ref}. In our experiments, $H_{low} = 150\ cm$, $H_{ref} = 30\ cm$,and $D_{ref} = 67\ pixels$. The background image is subtracted from the current frame (Figure 2(b)) to obtain the difference image (Figure 2(c)). A threshold is applied to the difference image to detect only those pixels having large intensity changes. Median filtering is applied to remove the salt and pepper noise in the difference image. The background subtraction is usually performed in color domain [12]. However, for the sake of simplicity in deciding the threshold value and in the median filtering operation, we performed the subtraction in the gray-scale domain. The difference image is scanned from the top to detect the top of the head and the vertical distance between the top of the head and the lowermost marker is measured as D_{user} (in pixels). An estimate of the true height of the person (H_{user} in cm) is computed as:

$$H_{user} = H_{low} + \frac{D_{user}}{D_{ref}} H_{ref}. \tag{1}$$

After the estimation of the height, the face of the user is detected in the captured frame using the algorithm proposed by Hsu et al. [13]. After the detection of the facial region in the frame (Figure 2(d)), the face is cropped out of the frame and is used by the face recognition and gender/ethnicity extraction modules. Since, we have not collected sufficient data using this extraction process, we used an off-line face database in our experiments.

3 Framework for Integration of Soft Biometrics

We use the same framework proposed in [14] for integrating the soft biometric information with the primary biometric system. In this framework, the biometric recognition system is divided into two subsystems. One subsystem is called the primary biometric system and it is based on traditional biometric identifiers like fingerprint, face and hand-geometry. The primary biometric system could be either unimodal or multimodal. The second subsystem, referred to as the secondary biometric system, is based on soft biometric traits like age, gender, and height. Figure 3 shows the architecture of a personal identification system that makes use of fingerprint, face and soft biometric measurements. Let $\omega_1, \omega_2, \cdots, \omega_n$ represent the n users enrolled in the database. Let

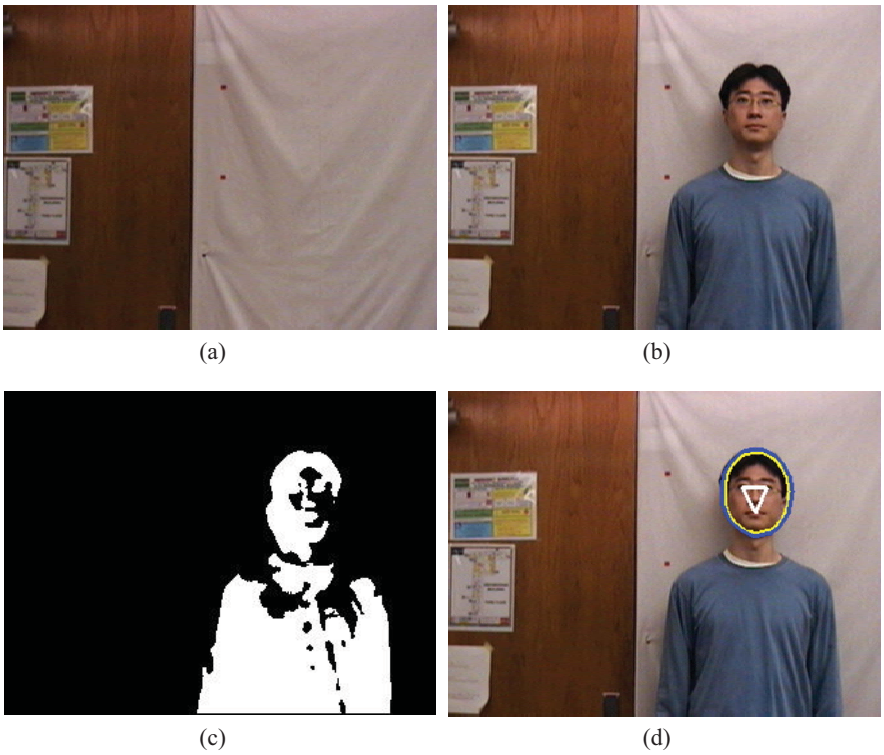

(a) (b)

(c) (d)

Fig. 2. Extraction of height and facial image from the user (a) background image (b) Current frame (c) Difference Image (d) Location of the face in the current frame.

x be the feature vector corresponding to the primary biometric. Without loss of generality, let us assume that the output of the primary biometric system is of the form $P(\omega_i \mid \mathbf{x}),\ i = 1, 2, \cdots, n$, where $P(\omega_i \mid \mathbf{x})$ is the probability that the test user is ω_i given the feature vector **x**. If the output of the primary biometric system is a matching score, it is converted into posteriori probability using an appropriate transformation. For the secondary biometric system, we can consider $P(\omega_i \mid \mathbf{x})$ as the prior probability of the test user being user ω_i.

Let $\mathbf{y} = [y_1,\ y_2,\ \cdots,\ y_k,\ y_{k+1},\ y_{k+2},\ \cdots,\ y_m]$ be the soft biometric feature vector, where y_1 through y_k are continuous variables and y_{k+1} through y_m are discrete variables. The updated probability of user ω_i, given the primary biometric feature vector **x** and the soft biometric feature vector **y** i.e., $P(\omega_i \mid \mathbf{x},\ \mathbf{y})$ can be calculated using the Bayes' rule.

$$P(\omega_i|\mathbf{x},\ \mathbf{y}) = \frac{p(\mathbf{y}|\omega_i)\ P(\omega_i|\mathbf{x})}{\sum_{i=1}^{n} p(\mathbf{y}|\omega_i)\ P(\omega_i|\mathbf{x})}\ . \tag{2}$$

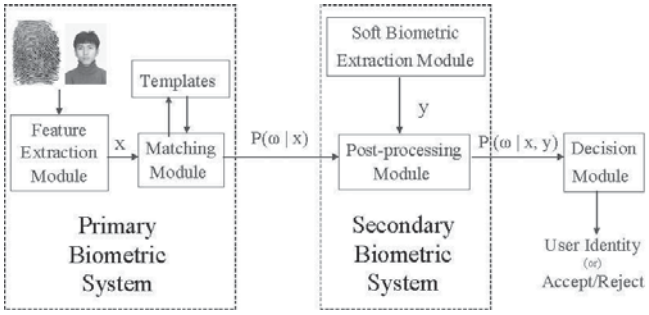

Fig. 3. Integration of Soft Biometric Traits with a Primary Biometric System (x is the fingerprint feature vector, y is the soft biometric feature vector).

If we assume that the soft biometric variables are independent, equation (2) can be rewritten as

$$P(\omega_i|\mathbf{x},\ \mathbf{y}) = \frac{p(y_1|\omega_i)\ \cdots\ p(y_k|\omega_i)\ P(y_{k+1}|\omega_i)\ \cdots\ P(y_m|\omega_i)\ P(\omega_i|\mathbf{x})}{\sum_{i=1}^{n} p(y_1|\omega_i)\ \cdots\ p(y_k|\omega_i)\ P(y_{k+1}|\omega_i)\ \cdots\ P(y_m|\omega_i)\ P(\omega_i|\mathbf{x})}\ . \tag{3}$$

In equation (3), $p(y_j|\omega_i),\ j = 1, 2, \cdots, k$ is evaluated from the conditional density of the variable y_j for user ω_i. On the other hand, discrete probability $P(y_j|\omega_i), j = k+1, k+2, \cdots, m$ represents the probability that user ω_i is assigned to the class y_j. This is a measure of the accuracy of the classification module in assigning user ω_i to one of the distinct classes based on biometric indicator y_j. In order to simplify the problem, let us assume that the classification module performs equally well on all the users and therefore the accuracy of the module is independent of the user.
Let

$$p(\mathbf{y}) = \sum_{i=1}^{n} p(y_1|\omega_i)\ \cdots\ p(y_k|\omega_i)\ P(y_{k+1}|\omega_i)\ \cdots\ P(y_m|\omega_i)\ P(\omega_i|\mathbf{x})\ .$$

The logarithm of $P(\omega_i|\mathbf{x},\ \mathbf{y})$ in equation (3) can be expressed as

$$\log P(\omega_i|\mathbf{x},\ \mathbf{y}) = \log p(y_1|\omega_i) + \ \cdots\ + \log p(y_k|\omega_i) + \log P(y_{k+1}|\omega_i) + \ \cdots$$
$$+ \log P(y_m|\omega_i) + \log P(\omega_i|\mathbf{x}) - \log p(\mathbf{y})\ . \tag{4}$$

This formulation has two main drawbacks. The first problem is that all the m soft biometric variables have been weighed equally. In practice, some variables may contain more information than the others. For example, the gender of a person may give more information about a person than height. Therefore, we must introduce a weighting

scheme for the soft biometric traits based on an index of distinctiveness and permanence; i.e., traits that have smaller variability and larger distinguishing capability will be given more weight in the computation of the final matching probabilities. Another potential pitfall is that any impostor can easily spoof the system because the soft characteristics have an equal say in the decision as the primary biometric trait. It is relatively easy to modify/hide one's soft biometric attributes by applying cosmetics and wearing other accessories (like mask, shoes with high heels, etc.). To avoid this problem, we assign smaller weights to the soft biometric traits compared to those assigned to the primary biometric traits. This differential weighting also has another implicit advantage. Even if a soft biometric trait of a user is measured incorrectly (e.g., a male user is identified as a female), there is only a small reduction in that user's posteriori probability and the user is not immediately rejected. In this case, if the primary biometric produces a good match, the user may still be accepted. Only if several soft biometric traits do not match, there is significant reduction in the posteriori probability and the user could be possibly rejected. If the devices that measure the soft biometric traits are reasonably accurate, such a situation has a low probability of occurrence. The introduction of the weighting scheme results in the following discriminant function for user ω_i:

$$g_i(\mathbf{x},\ \mathbf{y}) = a_0 \log P(\omega_i|\mathbf{x})\ +\ a_1 \log p(y_1|\omega_i) +\ \cdots\ +\ a_k \log p(y_k|\omega_i) + a_{k+1} \log P(y_{k+1}|\omega_i)\ +\ \cdots\ +\ a_m \log P(y_m|\omega_i), \quad (5)$$

where $\sum_{i=0}^{m} a_i = 1$ and $a_0 >> a_i,\ i = 1, 2, \cdots, m$. Note that a_i's, $i = 1, 2, \cdots, m$ are the weights assigned to the soft biometric traits and a_0 is the weight assigned to the primary biometric identifier. It must be noted that the weights $a_i,\ i = 1, 2, \cdots, m$ must be made small to prevent the domination of the primary biometric by the soft biometric traits. On the other hand, they must large enough so that the information content of the soft biometric traits is not lost. Hence, an optimum weighting scheme is required to maximize the performance gain.

4 Experimental Results

Our experiments demonstrate the benefits of utilizing the gender, ethnicity, and height information of the user in addition to the face and fingerprint biometric identifiers. The face database described in [8] has been used in our experiments. This database has face images of 263 users, with 10 images per user. Our fingerprint database consisted of impressions of 160 users obtained using a Veridicom sensor. Each user provided four impressions of each of the four fingers, namely, the left index finger, the left middle finger, the right index finger, and the right middle finger. Of these 640 fingers, 263 were selected and assigned uniquely to the users in the face database. A Linear Discriminant Analysis (LDA) based scheme is used for face matching. Eight face images of each user were used during the training phase and the remaining two images were used as test images. The face matching score vector (of length 263) was computed for each test image as follows. The similarity of the test image to the 2104 (263×8) training images in the database was found and the largest of the 8 scores of a particular user

was selected as the matching score for that user. Fingerprint matching was done using minutia features [15]. Two fingerprint impressions of each user were used as templates and the other two impressions were used for testing. The fingerprint matching score for a particular user was computed as the average of the scores obtained by matching the test impression against the two templates of that user. Thus, a fingerprint matching score vector for each test impression was computed. The separation of the face and fingerprint databases into training and test sets, was repeated 20 times and the results reported are the average for the 20 trials.

The ethnicity classifier proposed in [8] was used in our experiments. This classifier identifies the ethnicity of a test user as either Asian or non-Asian with an accuracy of 96.3%. If a "reject" option is introduced, the probability of making an incorrect classification is reduced to less than 1%, at the expense of rejecting 20% of the test images. A gender classifier was built following the same methodology used in [8] for ethnicity classification. The accuracy of the gender classifier without the "reject" option was 89.6% and the introduction of the "reject" option reduces the probability of making an incorrect classification to less than 2%. In cases where the ethnicity or the gender classifier cannot make a reliable decision, the corresponding information is not utilized for updating the matching score of the primary biometric system.

Since we did not have the height information about the users in the database, we randomly assigned a height 'H_i' to user ω_i, where the H_i's are drawn from a Gaussian distribution with mean 165 cm and a standard deviation of 15 cm. The height of a user measured during the recognition phase will not be equal to the true height of that user stored in the database due to the errors in measurement and the variation in the user's height over time. Therefore, it is reasonable to assume that the measured height H_i^* will follow a Gaussian distribution with a mean H_i cm and a standard deviation of 5 cm.

Let $P(\omega_i|s)$ be the posterior probability that the test user is user ω_i given the primary biometric score 's' of the test user. Let $y_i = (G_i,\ E_i,\ H_i)$ be the soft biometric feature vector corresponding to the user ω_i, where G_i, E_i, and H_i are the true values of gender, ethnicity, and height of ω_i. Let $y^* = (G^*,\ E^*,\ H^*)$ be the observed soft biometric feature vector of the test user, where G^* is the observed gender, E^* is the observed ethnicity, and H^* is the observed height. Now the final score after considering the observed soft biometric characteristics is computed as:

$$g_i(s,\ y^*) = a_0 \log P(\omega_i|s) + a_1 \log p(H^*|H_i) + a_2 \log P(G^*|G_i) + a_3 \log P(E^*|E_i)\ ,$$

where $a_2 = 0$ if G^* ="reject", and $a_3 = 0$ if E^* ="reject".

Experiments were conducted on three primary biometric systems, namely, fingerprint, face, and a multimodal system using face and fingerprint as the individual modalities. Figure 4 shows the Cumulative Match Characteristic (CMC) of the fingerprint biometric system operating in the identification mode, and the improvement in performance achieved after the utilization of soft biometric information. The weights assigned to the primary and soft biometric traits were selected experimentally such that the performance gain is maximized. However, no formal procedure was used and an exhaustive search of all possible sets of weights was not attempted. The use of ethnicity and gender information along with the fingerprint leads to an improvement of 1% in the rank one

performance as shown in Figures 4(a)and 4(b), respectively. From Figure 4(c), we can observe that the height information of the user is more discriminative than gender and ethnicity, and leads to a 2.5% improvement in the rank one performance. The combined use of all the three soft biometric traits results in an improvement of approximately 5% over the primary biometric system as shown in Figure 4(d).

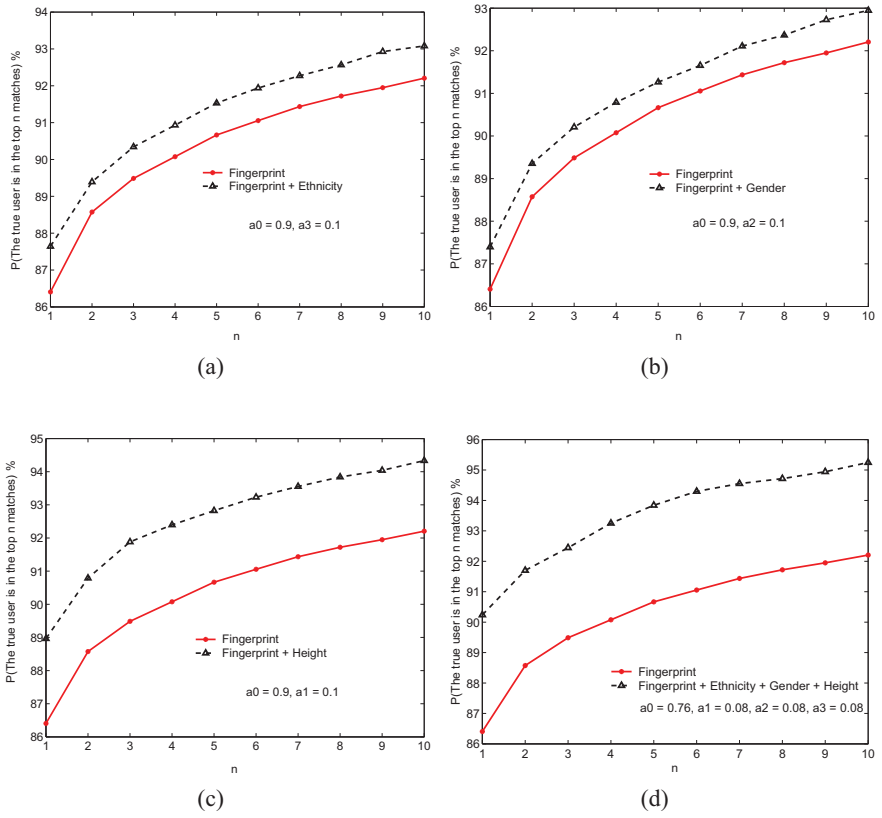

Fig. 4. Improvement in identification performance of fingerprint system after utilization of soft biometric traits.

The ethnicity and gender information did not provide any improvement in the performance of a face recognition system. This may be due to the fact that the gender and ethnicity classifiers, and the face recognition system use the same representation, namely, LDA for classification. The LDA algorithm for all the three classifiers operates on the same set of training images and hence it is highly likely that the features used for these classification problems are strongly correlated. However, the height information

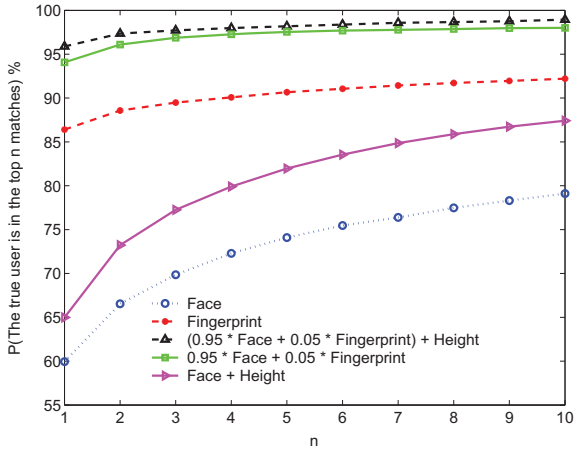

Fig. 5. Improvement in identification performance of (face + fingerprint) multimodal system after utilization of the height of the user.

is independent of the facial features and, hence, it leads to an improvement of 5% in the face recognition performance (see Figure 5). The failure of the ethnicity and gender information to improve the face recognition performance establishes that fact that soft biometric traits would help in recognition only if the identity information provided by them is complementary to that of the primary biometric identifier.

Figure 5 shows the CMC curves for a multimodal system using face and fingerprint as the individual modalities. In this system, the combined matching score of the primary biometric system is computed as a weighted average of the scores of the face and fingerprint modalities. We can observe that the rank one performance of this multimodal system is superior to that of the individual modalities by 8%. The addition of height as a soft biometric feature further improves the performance by 2%. This shows soft biometric traits can be useful even if the primary biometric system already has a high accuracy.

5 Summary and Future Directions

We have demonstrated that the utilization of ancillary user information like gender, height, and ethnicity can improve the performance of the traditional biometric systems. Although the soft biometric characteristics are not as permanent and reliable as the traditional biometric identifiers like fingerprint, they provide some information about the identity of the user that leads to higher accuracy in establishing the user identity. We have also shown that soft biometric characteristics would help only if they are complementary to the primary biometric traits. However, an optimum weighting scheme based the discriminative abilities of the primary and the soft biometric traits is needed to achieve an improvement in recognition performance.

Our future research work will involve establishing a more formal procedure to determine the optimal set of weights for the soft biometric characteristics based on their

distinctiveness and permanence. Methods to incorporate time-varying soft biometric information such as age and weight into the soft biometric framework will be studied. The effectiveness of utilizing the soft biometric information for "indexing" and "filtering" of large biometric databases must be studied. Finally, more accurate mechanisms must be developed for automatic extraction of soft biometric traits.

References

1. Jain, A.K., Bolle, R., Pankanti, S., eds.: Biometrics: Personal Identification in Networked Security. Kluwer Academic Publishers (1999)
2. Hong, L., Jain, A.K., Pankanti, S.: Can Multibiometrics Improve Performance? In: Proceedings of IEEE Workshop on Automatic Identification Advanced Technologies, New Jersey, U.S.A. (1999) 59–64
3. Heckathorn, D.D., Broadhead, R.S., Sergeyev, B.: A Methodology for Reducing Respondent Duplication and Impersonation in Samples of Hidden Populations. In: Annual Meeting of the American Sociological Association, Toronto, Canada (1997)
4. Gutta, S., Huang, J.R.J., Jonathon, P., Wechsler, H.: Mixture of Experts for Classification of Gender, Ethnic Origin, and Pose of Human Faces. IEEE Transactions on Neural Networks **11** (2000) 948–960
5. Moghaddam, B., Yang, M.H.: Learning Gender with Support Faces. IEEE Transactions on Pattern Analysis and Machine Intelligence **24** (2002) 707–711
6. Shakhnarovich, G., Viola, P., Moghaddam, B.: A Unified Learning Framework for Real Time Face Detection and Classification. In: Proceedings of International Conference on Automatic Face and Gesture Recognition, Washington D.C., USA (2002)
7. Balci, K., Atalay, V.: PCA for Gender Estimation: Which Eigenvectors Contribute? In: Proceedings of Sixteenth International Conference on Pattern Recognition. Volume 3., Quebec City, Canada (2002) 363–366
8. Jain, A.K., Lu, X.: Ethnicity Identification from Face Images. In: Proceedings of SPIE International Symposium on Defense and Security : Biometric Technology for Human Identification (To appear). (2004)
9. Woodward, J.D.: Testimony to the Commission on Online Child Protection on Age Verification Technologies. available at `http://www.copacommission.org/meetings/hearing1/woodward.test.pdf` (2000)
10. Buchanan, M.V., Asano, K., Bohanon, A.: Chemical Characterization of Fingerprints from Adults and Children. In: Proceedings of SPIE Photonics East Conference. Volume 2941. (1996) 89–95
11. Kwon, Y.H., Lobo, N.V.: Age Classification from Facial Images. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition. (1994) 762–767
12. Hong, D., Woo, W.: A Background Subtraction for a Vision-based User Interface. In: Proceedings of Fourth International Conference on Information, Communications and Signal Processing, Pacific-Rim Conference On Multimedia, 1B3.3. (2003)
13. Hsu, R.L., Mottaleb, M.A., Jain, A.K.: Face Detection in Color Images. IEEE Transactions on Pattern Analysis and Machine Intelligence **24** (2002) 696–706
14. Jain, A.K., Dass, S.C., Nandakumar, K.: Can soft biometric traits assist user recognition? In: Proceedings of SPIE International Symposium on Defense and Security : Biometric Technology for Human Identification (To appear). (2004)
15. Jain, A.K., Hong, L., Pankanti, S., Bolle, R.: An identity authentication system using fingerprints. Proceedings of the IEEE **85** (1997) 1365–1388