

Gender Classification from Hand Shape

Gholamreza Amayeh, George Bebis, Mircea Nicolescu

Computer Vision Laboratory, University of Nevada, Reno, NV 89557

{amayeh, bebis, mircea}@cse.unr.edu

Abstract

Many social interactions and services today depend on gender. In this paper, we investigate the problem of gender classification from hand shape. Our work has been motivated by studies in anthropometry and psychology suggesting that it is possible to distinguish between male and female hands by considering certain geometric features. Our system segments the hand silhouette into six different parts corresponding to the palm and fingers. To represent the geometry of each part, we use region and boundary features based on Zernike moments and Fourier descriptors. For classification, we compute the distance of a given part from two different eigenspaces, one corresponding to the male class and the other corresponding to female class. We have experimented using each part of the hand separately as well as fusing information from different parts of the hand. Using a small database containing 20 males and 20 females, we report classification results close to 98% using score-level fusion and LDA.

1. Introduction

Gender classification is an important problem with a variety of practical applications. For example, a robust gender classification system could provide a basis for performing passive surveillance using demographic information or collecting valuable consumer statistics in a shopping center. It could also be used to improve the performance of biometric systems such as face authentication and recognition [16]. In computer vision, the majority of studies on gender classification are based on face since visual information from human faces provides important cues for gender classification. A recent study comparing different gender classification approaches using face information can be found in [10]. A very small number of studies have also investigated the use of modalities other than face including gait, [13], iris [15] and fingerprint [4].

In this study, we investigate the problem of gender classification from human hands. To the best of our knowledge, this is the first study attempting to address the issue of

gender classification from hand images in computer vision. However, extracting gender information from human hands has been studied for a long time fields such as anthropology and psychology. For example, several studies have found that there is significant difference between the hand dimensions of males and females. In 1875, Ecker [7] noted that the following three relations of relative finger length may be observed in a human hand: *i*) the index finger is shorter than the ring finger; *ii*) the index finger is equal in length to the ring finger; and *iii*) the index finger is longer than ring finger. In 1877, Mantegazza [11] found that all three relations occur in both sexes, however, a relatively long index finger is found more frequently in females than in males. In 1930, George [8] studied a Canadian population and found that relation (*i*) is more frequent in males while relation (*ii*) is predominant in females. Similar results have been reported in more recent studies [5].

In [12], McFadden and Shubel studied all six possible ratios between the index, middle, ring and little fingers in both genders. Their results indicate that the ratio exhibiting the largest gender difference was the relative lengths of the index and ring fingers. In another study, Agnihotri et al. [1] found that the average hand breadth and hand length are about 1 cm and 1.5 cm correspondingly greater in male subjects than in female subjects. By defining hand index as the ratio of hand breadth over hand length, they found that the average hand index in males was more than 44% while the average index in females was less than 44%. Table 1 provides more details. Based on these results, they suggested using this value as a threshold for determining gender by hand dimensions. However, to date there is no conclusive evidence as to which features determines gender robustly. It appears that gender can not be determined using a single feature, but rather involves the combination of multiple features.

Motivated by these studies, we investigate the problem of gender classification from hand images by extracting more powerful hand features. Our goal is building a system that can distinguish between male and female subjects using hand shape information. Our system segments the hand silhouette into six different parts corresponding to the

Table 1. Measurements of hand breadth and length (based on centimeters) [1].

	Population	Hand Breadth (cm)				Hand Length (cm)				Hand Index (%)
Gender	N	min	max	μ	σ	min	max	μ	σ	μ
Male	125	7.3	9.4	8.45	0.40	15.3	21.0	18.89	0.88	44.73
Female	125	6.7	8.8	7.48	0.38	14.8	20.4	17.22	0.92	43.46

palm and fingers using a methodology proposed in our recent work on hand-based authentication [2][3]. To represent the geometry of each part, we use region and boundary features based on Zernike moments (ZMs) [9] and Fourier descriptors (FDs) [17]. For classification, we compute the distance of a given part from two different eigenspaces, one corresponding to the male class and the other corresponding to female class. We have experimented using each part of the hand separately as well as fusing information from different parts of the hand.

Determining gender information from hand images has several important advantages. First of all, capturing an image of the hand can be done more robustly than capturing an image of the face. There are several biometric systems today that can capture high quality hand images by controlling the position and orientation of the hand as well as illumination [14]. Second, assuming that the hand is placed on a flat surface for image acquisition purposes, which is typical for hand-based authentication applications, hand appearance shows less variability compared to face appearance which is affected by many factors such as facial expression changes. Finally, gender information from hand images could be very valuable in improving the accuracy and robustness of hand-based authentication and identification systems [2][3].

2. Image Acquisition and Preprocessing

For image acquisition, we have used the system reported in [2] for hand-based authentication. This system consists of a CCD camera and a flat lighting table, which forms the surface for placing the hand. The direction of the camera is perpendicular to the lighting table. The camera has been calibrated to remove lens distortion. The size of the captured images is 480×640 pixels. Figures 1(a) and (b) show two sample images acquired by this system. Each image goes through a preprocessing stage. First, the image is binarized using thresholding. Since image quality is very high due to our set-up (i.e., almost free of shadows and noise), a fixed threshold works well. After binarization, we segment the hand silhouette into six different regions corresponding to the palm and the fingers. Segmentation is performed using an iterative process based on morphological filters [3]. Figure 1(c) shows the segmentation results on the image of Figure 1(a).




Figure 1. (a) Female hand image, (b) male hand image, and (c) segmented female hand image.

3. Feature Extraction

Once the hand silhouette has been segmented into different regions, each region is represented by a set of features. We have experimented with two different MPG-7 shape descriptors [19] to represent the geometry of the fingers and the palm. In particular, MPG-7 divides shape descriptors in two categories: *contour-based* and *region-based*. Contour-based shape descriptors use the shape's boundary to extract shape information, while region-based shape descriptors exploit the shape's region to represent shape information. In this work, we experimented with ZMs and FDs, both of which are listed in the MPEG-7 standard.

3.1. Zernike Moments

ZMs are based on a set of complex polynomials that form a complete orthogonal set over the interior of the unit circle [9]. They are defined as the projection of the image on these orthogonal basis functions. Specifically, the basis functions $V_{n,m}(x, y)$ are given by

$$V_{n,m}(x, y) = V_{n,m}(\rho, \theta) = R_{n,m}(\rho)e^{jm\theta} \quad (1)$$

where n is a non-negative integer, m is a non-zero integer subject to the constraints $n - |m|$ is even and $|m| < n$, ρ is the length of the vector from origin to (x, y) , θ is the angle between the vector ρ and the x -axis in a counter clockwise direction, and $R_{n,m}(\rho)$ is the Zernike radial polynomial. $R_{n,m}(\rho)$ is defined as follows:

$$\begin{aligned} R_{n,m}(\rho) &= \sum_{k=|m|, n-k=\text{even}}^n \frac{(-1)^{\frac{n-k}{2}} \frac{n+k!}{\frac{n-k}{2}! \frac{k+m}{2}! \frac{k-m}{2}!}}{\rho^k} \\ &= \sum_{k=|m|, n-k=\text{even}}^n \beta_{n,m,k} \rho^k \end{aligned} \quad (2)$$

The Zernike moment of order n with repetition m for a digital image function $f(x, y)$ is given by [9]:

$$Z_{n,m} = \frac{n+1}{\pi} \sum_{x^2+y^2 \leq 1} \sum_{y^2 \leq 1-x^2} f(x,y) V_{n,m}^*(x,y) \quad (3)$$

where $V_{n,m}^*(x,y)$ is the complex conjugate of $V_{n,m}(x,y)$. To compute the Zernike moments of a given image, the image center of mass is taken to be the origin.

When using ZMs, one has to deal with several practical issues including computational cost of high-order ZMs and lack of accuracy due to limited numerical precision. In this work, we employ an improved algorithm, proposed in one of our earlier works [2][3], which speeds up computations by exploiting some recursive relations in the computation of the ZMs and by using look-up tables. A crucial parameter here is determining the maximum ZM order to represent the geometry of different parts of the hand. Amayeh et al. [3] found that using ZMs up to order 20 and 30 for the fingers and palm, respectively, captures enough information for reconstructing them precisely. Therefore, we used the same parameters here.

3.2. Fourier Descriptors

FDs have been used in a wide range of applications to describe the boundary of an object. Let us consider a closed contour C in the complex plane. In this case, the x - y coordinates of each point in the boundary become a complex number $x + jy$. By tracing the boundary in a counterclockwise direction with uniform velocity, a complex function $z(t) = x(t) + jy(t)$ is obtained with parameter t . The velocity is chosen such that the time required to traverse the contour is 2π . If $z(k)$ is a uniformly re-sampled version of $z(t)$ of dimension N , its Discrete Fourier Transform (DFT) is given by following equation:

$$z(k) = \sum_{n=0}^N a_n e^{\frac{j2\pi nk}{N}} \quad (4)$$

where a_n is Fourier coefficient of $z(k)$. The FDs of the closed contour C are defined by taking the inverse transform:

$$a_n = \frac{1}{N} \sum_{k=0}^N z(k) e^{-\frac{j2\pi nk}{N}}, \quad n \in \{1, 2, \dots, N\} \quad (5)$$

To normalize the FDs with respect to translation, rotation, scale, and starting point, we used the methodology proposed in [17]. In [17], the dimensionality of the contour (N) must be power of 2. Since the average number of points in the boundary of different parts of the hand was in the range of $[2^7, 2^8]$ for the fingers and $[2^8, 2^9]$ for the palm, we re-sample the finger and palm contours in 256 and 512 points respectively.

4. Gender Classification Using Different Parts of the Hand

To investigate the discrimination power of different parts of the hand for gender classification, we considered each part of the hand separately. To represent each part compactly, we applied Principal Component Analysis (PCA) [6]. Specifically, we built two different eigenspaces for each part of the hand, one for the male class and the other for the female class. Given that the hand is decomposed in six parts (i.e., five fingers and the palm), we built a total of twelve eigenspaces.

To represent an instance of given part, we compute its distance from the male and female eigenspaces. This was performed by projecting the part in each eigenspace and reconstructing it from its projections. To compute the error in each eigenspace, we computed the difference between the original representation of the part and its reconstruction. Specifically, let us assume that $\Omega_{m/f}$ corresponds to the representation of an instance Φ of some part p in the male/female eigenspaces; then $\Omega_{m/f}$ is given by:

$$\Omega_{m/f} = \sum_{k=0}^M \omega_{m/f}^k u_{m/f}^k + \bar{\Phi}_{m/f} \quad (6)$$

where the projection $\omega_{m/f}^k$ of Φ in the male/female eigenspaces can be computed as follows:

$$\omega_{m/f}^k = \hat{u}_{m/f}^k (\Phi - \bar{\Phi}_{m/f}) \quad (7)$$

M represents the dimensionality of the eigenspaces, $u_{m/f}^k$ is the k th eigenvector in the male/female space and $\bar{\Phi}_{m/f}$ is the average male/female vector (i.e., computed from the training set). Also, $\hat{u}_{m/f}^k$ is the transpose of $u_{m/f}^k$. To measure the *masculine/feminine* characteristic of Φ , the Euclidean distance $\varepsilon_{m/f}$ between Φ and its projection onto the male/female eigenspaces is computed:

$$\varepsilon_{m/f} = \|\Phi - \Omega_{m/f}\| \quad (8)$$

Therefore, each part p is represented as a distance vector $E = [\varepsilon_m, \varepsilon_f]^T$. Figure 2 illustrates this process. In our experiments, we preserved the same amount of information for the male/female eigenspaces, however, this could be varied depending on the shape descriptor.

Figure 3 shows the distribution of distances $\varepsilon_{m/f}$ in the case of the little finger using ZMs and FDs. Due to lack of space, we do not show the distributions of the other parts of the hand. This male/female eigenspaces in this figure were generated using 12 males and 12 females from our database and preserving 99% information.

To classify a query distance vector E , we experimented with three different classifiers: *Minimum Distance* (MD), *k-Nearest Neighbors* (kNN), and *Linear Discriminant Analysis* (LDA) [6]. In the case of the MD classifier, the minimum




Figure 2. Computing the masculine/feminine characteristic for an instance Φ of some part of the hand.

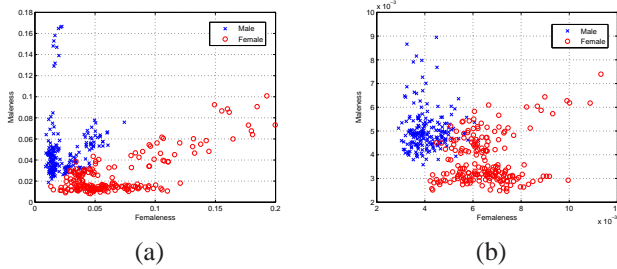


Figure 3. Distribution of distances for the male/female eigenspaces in the case of little finger: (a) ZMs, (b) FDs.

distance from the male/female eigenspace was used to determine the gender of the input part. In the case of a tie, we arbitrarily classified the input as male. In the case of the kNN classifier, we determined the gender by finding the most dominant gender classification among the top k closest distances between the query and the training data. Again, in the case of a tie, we arbitrarily classified the input as male. In the case of LDA, each distance vector was represented by a single value since our problem is a two-class classification problem. A threshold needs to be used in this case in order to separate the two classes. Figure 4 shows the male/female distributions for the little finger using LDA. These graphs should be related to the classification results presented in Section 6.

5. Gender Classification Using Fusion

Different parts of the hand have their inherent strengths and weakness. Fusing information from the fingers and palm has the potential to improve overall classification performance. Recently, Amayeh et al. [3] have considered the problem of fusing information from the fingers and the palm for hand-based authentication, illustrating accuracy improvements. Similarly, we propose combining information from the fingers and the palm in order to improve gender classification accuracy and robustness. We have ex-

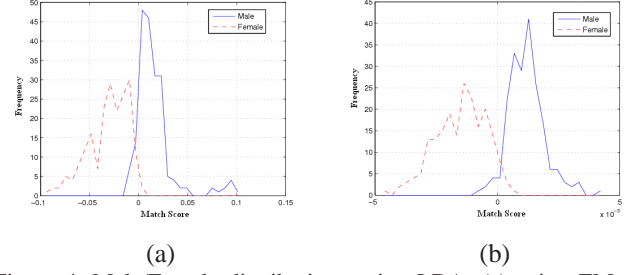


Figure 4. Male/Female distributions using LDA: (a) using ZMs distances shown in Figure 3(a), (b) using FD distances shown in Figure 3(b).

perimented with three different fusion strategies: *feature-level fusion*, *score-level fusion* and *decision-level fusion*. It should be mentioned that we have also experimented with extracting boundary and region features from the whole hand (i.e., without segmenting the hand in different parts), however, finger motion causes many problems. Segmenting the hand in different parts and extracting features from each part separately addresses the issue of finger motion effectively.

5.1. Feature-Level Fusion

For feature-level fusion, we concatenate the features from each part of the hand into one vector and apply PCA to reduce its dimensionality. Then, we build a separate eigenspace for each class (i.e., male/female) and perform classification as described in Section 4. For comparison purposes, we have experimented with fusing the original feature vectors for each part, based on ZMs and FDs, as well as the distance vectors E .

5.2. Score-Level Fusion

In score-level fusion, the matching scores from the fingers and the palm are fused together into an overall score. To compute the matching score for each part, we used the output of LDA (see Section 4). The weighted sum rule is the most straightforward fusion strategy at the score level. In this case, the matching scores are combined into a single score by applying an appropriate weight on each matching score as follows:

$$S = \sum_{i=1}^6 w_i s_i \quad (9)$$

where s_i is the score of the i -th part and w_i is the weight associated with this part; the weights must satisfy $\sum w_i = 1$. The first five terms of the sum correspond to the little, ring, middle, index and thumb fingers respectively, while the last term corresponds to the palm.

5.3. Decision-Level Fusion

In decision-level fusion, the final classification is obtained by fusing independent classifications based on different parts of the hand. To classify each part of the hand, we use LDA (see Section 4). To fuse the individual classifications, we use *majority voting* which is among the most straightforward decision-level fusion strategies. First, the gender of each part of the hand is determined using LDA. Then, the subject is classified as male if three or more parts of the subject’s hand have been classified as a male; otherwise, the subject is classified as female.

6. Experimental Results

To evaluate the proposed approach, we built a small database containing 40 people. Although we have a larger database from our previous work on hand-based authentication [3], unfortunately, we have not recorded gender information for each subject in that database. The population of male and female subjects was equal (i.e., 20 males and 20 females). For each subject, we collected 10 images of his/her right hand in different directions and positions. Time lapse between captured samples of same subject was a few minutes. Besides asking the subjects to stretch their hand and place it inside a square area drawn on the surface of the lighting table, no other restrictions were imposed.

To evaluate the proposed system, we used cross-validation based on *leave-one-out* approach. Although leave-one-out is computationally expensive, it is well suited for small datasets. Moreover, it has been shown to be an almost unbiased estimator of the true error rate of a classifier [18]. Using a leave-one-out approach, we repeated each experiment 40 times, each time using all samples of one person (10 out of 400 samples) for test and the rest of samples (390 out of 400 samples) for training set. Therefore, for each experiment, we report the average over 40 trials. Zernike moments were computed up to order 20 and 30 for the fingers and palm respectively. The number of features was 121 for each finger and 256 for the palm. The number of Fourier coefficients was 256 for the fingers and 512 for the palm.

6.1. Gender Classification Results using Different Parts of the Hand

In this subsection, the performance of each part of the hand is presented using different shape descriptors and classifiers as described in Section 4. Different portion of information has been preserved in building the male/female eigenspaces (i.e. 90%, 95%, 97% and 99%).

Figures 5 - 10 show gender classification results using different parts of the hand and different combinations of features (i.e., ZMs or FDs) and classifiers (i.e., MD, kNN, and LDA). In particular, figures 5, 6 and 7 show gender




Figure 5. Gender classification results using ZMs and MD classifier.

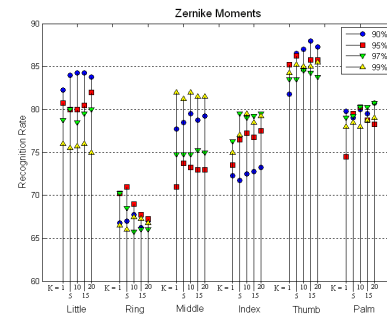


Figure 6. Gender classification results using ZMs and k -NN classifier where $k \in \{1, 5, 10, 15, 20\}$.

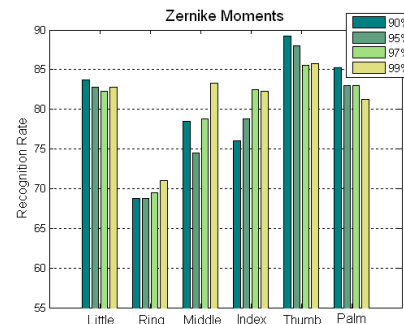


Figure 7. Gender classification results using ZMs and LDA.

classification results using Zernike moment. The results are consistent across all classifiers, showing that the best accuracy is achieved by the thumb while the lowest accuracy is achieved by the ring finger. When considering FDs (i.e., Figures 8, 9 and 10), the ring finger shows comparable performance with the rest of parts.

Comparing ZMs with FDs, the results are rather mixed. For certain parts (e.g., ring finger), FDs seem to be performing better, while for others (thumb and palm), ZMs seem to be performing better. An interesting observation, however, is that the amount of information preserved when building the male/female eigenspaces seems to affect the FDs more than the ZMs. This is illustrated more clearly in Figures

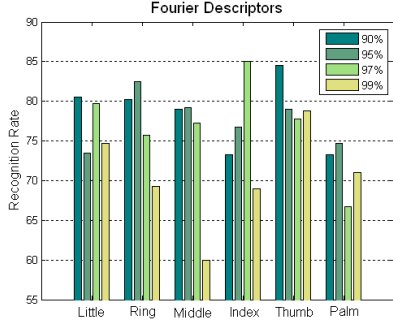


Figure 8. Gender classification results using FDs and MD classifier.

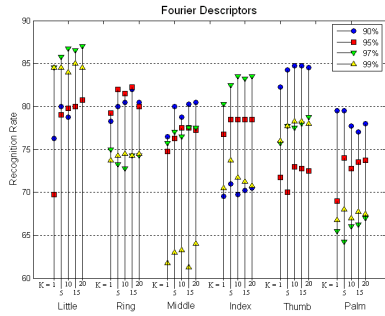


Figure 9. Gender classification results using FDs and k -NN classifier where $k \in \{1, 5, 10, 15, 20\}$.




Figure 10. Gender classification results using FDs and LDA.

6 and 9. Overall, the highest accuracy in the case of ZMs was obtained using the thumb and the LDA classifier (i.e., close to 89%). In the case of FDs, the highest accuracy was obtained using the little finger and either the kNN or LDA classifiers (i.e., close to 87%).

6.2. Gender Classification Results using Fusion

In this section, we report classification results by fusing information from different parts of the hand. Since LDA performed the best in our previous experiments, we only considered LDA for classifying different parts of the hand for fusion purposes. Table 2 shows the classification results obtained using feature-level fusion. As discussed in Sec-

Table 2. Gender classification using feature-level fusion.

Preserved info.	Zernike Moments		Fourier Descriptor	
	Concat. Feat. Vec.	Concat. Dist. Vec.	Concat. Feat. Vec.	Concat. Dist. Vec.
90% (PCA)	91.75	93.5	86.75	92
95% (PCA)	90.75	92.25	88	88.25
97% (PCA)	92.25	91	88.75	87.75
99% (PCA)	93	89	88.5	89.75

Table 3. Gender classification using score-level fusion.

Preserved info.	Zernike Moments		Fourier Descriptor	
	Sum Rule	Weighted Sum	Sum Rule	Weighted Sum
90% (PCA)	94.5	97.75	95.25	98
95% (PCA)	91.25	96	92.75	96
97% (PCA)	91.75	96.25	91.25	96
99% (PCA)	92.25	96.5	92.25	95.75

tion 5, we experimented with fusing two types of features: (i) the original ZMs or FDs and (ii) the distance vectors E . As it can be observed, feature-level fusion improved classification results in certain cases, both for ZMs and FDs.

Table 3 shows the classification results obtained using score-level fusion (i.e., weighted sum rule). The main issue with this method is determining a set of appropriate weight values. Here, we chose the weight values by considering the discrimination power of each part of the hand as described in Section 4. For comparison purposes, we also experimented with the simple *Sum rule* where all the weights are equal (i.e., average over scores). As it can be observed, score-level fusion improves classification results significantly, both for ZMs and FDs. As expected, the weighted sum rule outperforms the simple sum rule.

Table 4 shows the classification results obtained using decision-level fusion (i.e., majority voting). As it can be observed, decision-level fusion improves classification results both for ZMs and FDs. Figure 11 shows the best classification rates obtained using different fusion strategies and features. Comparing all three fusion strategies, it is clear that score-level fusion outperforms the other two. Since different parts of the hand have different discrimination power, it is not surprising that score-level fusion using weighted-sum outperforms decision-level fusion. Obviously, weighting each part of the hand appropriately has an important impact on system error. It is worth noticing that when each part of the hand is assigned the same weight (i.e., sum rule), score-level fusion performs almost the same as decision-level fusion. Overall, the best performance (i.e., 98%) was obtained using score-level fusion and FDS. It should be observed, however, that the performance of ZMs using score-level fusion is very close to that of FDs.

Table 4. Gender classification using decision-level fusion.

Majority Voting	Preserved information			
	90% (PCA)	95% (PCA)	97% (PCA)	99% (PCA)
Zernike Moments	94.5	91.25	91.75	92.25
Fourier Descriptors	95.25	92.75	91.25	92.25

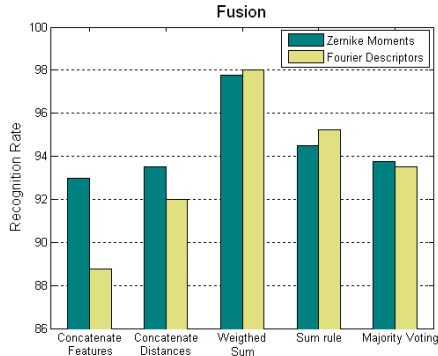


Figure 11. Best gender classification results using different fusion strategies and shape descriptors.

7. Conclusion

We investigated the problem of gender classification from hand images. The proposed system decomposes the hand silhouette in different parts corresponding to the fingers and the palm and describes the geometry of each part using either boundary descriptors, based on FDs, or region descriptors, based on ZMs. To classify a given hand as male or female, it fuses information from different parts of the hand using score-level fusion and LDA. Although the dataset used in our experiments is rather small and all the data was obtained in the same session, we believe that the results presented in this study are quite encouraging.

For future work, first we plan to perform more experiments using larger datasets. Moreover, we plan to investigate the issue of fusing boundary and region descriptors to further improve performance. Our experimental results in Section 6.1 indicate that ZMs and FDs provide complementary information. Second, we plan to evaluate the effect of time lapse on system performance. Third, we plan to evaluate the proposed system on populations of different ages. For example, an interesting problem would be to investigate whether hand shape can be used to distinguish the gender of children. Finally, we plan to investigate the benefits of integrating hand-based gender classification with hand-based authentication/identification.

References

[1] A. K. Agnihotri, B. Purwar, N. Jeebun, and S. Agnihotri. Determination of sex by hand dimensions. *The Internet Journal of Forensic Science*, 1(2), 2006.

[2] G. Amayeh, G. Bebis, A. Erol, and M. Nicolescu. Peg-free hand shape verification using high order zernike moments. *IEEE Computer Society Workshop on Multi-modal Biometrics*, June 2006.

[3] G. Amayeh, G. Bebis, A. Erol, and M. Nicolescu. A component-based approach to hand verification. *IEEE Computer Society Workshop on Biometrics*, June 2007.

[4] A. Badawi, M. Mahfouz, R. Tadross, and R. Jantz. Fingerprint-based gender classification. *The International Conference on Image Processing, Computer Vision, and Pattern Recognition*, June 2006.

[5] W. Brown, M. Hines, B. Fane, and M. Breedlove. Masculinized finger length patterns in human males and females with congenital adrenal hyperplasia. In *Horm Behavior*, volume 42(4), pages 380–385, 2002.

[6] R. Duda, P. Hart, and D. Stork. *Pattern Classification*. Wiley-Interscience, 2nd edition, 2000.

[7] A. Ecker. Einige bemerkungen tiber ein schwankenden charakter in der hand des menschen. *Arch. Anthropol. Brunschw.*, 8:67–74, 1875.

[8] R. George. Human finger types. *Anatomical Record*, 46:199–204, 1930.

[9] A. Khotanzad and Y. Hong. Invariant image recognition by zernike moments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:489–498, 1990.

[10] E. Makinen and R. Raisamo. Evaluation of gender classification methods with automatically detected and aligned faces. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 30(3), pages 541–547, 2008.

[11] P. Mantegazza. Della lunghezza relativa dell'indice a dell'anulare nella mano umana. *Arch. Anthropol. Etnol.*, 7(1):19–25, 1877.

[12] D. McFadden and E. Shubel. Relative lengths of fingers and toes in human males and females. *Hormones and Behavior*, 42:492–500, 2002.

[13] C. Shan, S. Gong, and P. W. McOwan. Learning gender from human gaits and faces. *IEEE Conference on Advanced Video and Signal Based Surveillance*, pages 505–510, September 2007.

[14] D. Sidlauskas and S. Tamer. Hand geometry recognition. In *Handbook on Biometrics (A. Jain, P. Flynn and A. Ross eds)*. Springer-Verlag, 2008.

[15] V. Thomas, N. V. Chawla, K. W. Bowyer, and P. J. Flynn. Learning to predict gender from iris images. *IEEE International Conference on Biometrics: Theory, Applications, and Systems*, pages 1–5, September 2007.

[16] K. Veropoulos, G. Bebis, and M. Webster. Investigating the impact of face categorization on recognition performance. In *International Symposium on Visual Computing (LNCS, vol 3804)*, December, 2005.

[17] T. Wallace and P. Wintz. An efficient three-dimensional aircraft recognition algorithm using normalized fourier descriptors. *Computer Vision Graphics Image Processing*, 13, 1980.

[18] S. Weiss and C. Kulikowski. *Computer Systems that Learn*. Morgan Kaufmann, 1991.

[19] D. Zhang and G. Lu. Evaluation of mpeg-7 shape descriptors against other shape descriptors. In *Multimedia Systems*, volume 9(1), pages 15–30, 2003.