FACE RECOGNITION WITH ENHANCED PRIVACY PROTECTION

Yongjin Wang, Dimitrios Hatzinakos

University of Toronto
The Edward S. Rogers Sr. Department of Electrical and Computer Engineering
10 King's College Road, Toronto, ON, Canada, M5S 3G4

ABSTRACT

This paper presents a novel approach for face based biometric recognition. The proposed method is based on the sorted index numbers (SIN) of appearance based facial features. A new algorithm is introduced to measure the similarity between SIN vectors. Due to the non-invertibility of the transformation from the original features to the SIN vectors, the proposed method can preserve the privacy of the users. The effectiveness of the proposed method is tested on a large generic data set, which contains images from several well known face databases. Experimental results demonstrate that the proposed solution may improve the recognition accuracy in both identification and verification scenarios.

Index Terms— Face recognition, Privacy, Index numbers

1. INTRODUCTION

Biometric recognition has been an active research area in the past two decades. Biometrics based recognition systems determine or confirm the identity of an individual based on the physiological and/or behavioral characteristics. A wide variety of biometric modalities have been investigated in the past. Examples of these biometrics include physiological traits such as fingerprint, face, iris, and behavioral characteristics such as gait and keystroke. Depending on different application context, a biometric system can operate in identification mode or verification mode. Biometric identification is a one-to-many comparison to find an individual's identity. A biometric verification system is a one-to-one match that determines whether the claim of an individual is true.

While biometric technology provides various advantages, there exist some problems. In the first place, biometrics reflect the user's physiological/behavior characteristics. The user's privacy may be compromised if the biometric templates is obtained by an adversary. The biometric templates should be stored in a format such that the user's privacy is preserved even the storage device is compromised. Secondly, biometrics can not be easily changed and reissued if compromised due to the limited number of biometric traits that human has. This is particularly important in biometric verification scenarios. Ideally, just like password, the biometric templates should be changeable. The users may use different biometric representation for different applications. When the biometric template in one application is compromised, the biometric signal itself is not lost forever and a new biometric template can be issued.

Yongjin Wang would like to acknowledge the Natural Sciences and Engineering Research Council of Canada (NSERC) for financial support.

A number of research works have been proposed in recent years to address the changeability and privacy preserving problems of biometric systems. One approach is to combine biometric technology with cryptographic systems [1]. In a biometric crypto-system, a randomly generated cryptographic key is combined with the biometric features in a secure way, and error correction algorithms are usually employed to tolerant errors. Due to the binary nature of the keys, such systems usually require discrete representation of biometric data. Representative works of this method include the fuzzy commitment scheme [2], fuzzy vault [3], fuzzy extractor [4], and helper data system [5]. An alternative solution is to apply repeatable and non-invertible transformations on the biometric features [6]. With this method, every enrollment (or application) can use a different transform. When a biometric template is compromised, a new one can be generated using a new transform. The major challenge here lies in the difficulty of preserving the similarity measure in the transformed domain. Existing works following this line include the BioHashing technique in [7] and convolution based method in [8]. However, although advances have been achieved in the past, existing works either can not provide robust privacy protection [1][9], or sacrifice recognition accuracy for privacy protection.

Among various biometrics, face recognition has been one of the most passive, natural, and noninvasive types of biometrics. Such characteristics of face recognition make it a good choice for a wide variety of applications such as surveillance, physical access control, computer network login, and ATM. Many face recognition methods have been proposed in the literature, among which appearance based approaches (such as principal component analysis and linear discriminant analysis) that treat the face image as a holistic pattern seem to be the most successful [10]. In this paper, we introduce a novel method for privacy preserving face recognition based on appearance based facial features. Unlike traditional face recognition methods which store either the original image or facial features as templates, the proposed method stores the sorted index numbers only. A matching algorithm is introduced to measure the similarity between two vectors of sorted index numbers. Because it is impossible to recover the original features based on the index numbers, the privacy of the users can be protected. As it will be shown, the proposed method can also be combined with cryptographic techniques as well as intentional transformation methods for changeable biometric template generation.

The remainder of this paper is organized as follows. Section 2 introduce the proposed method. Experimental results along with detailed discussion are presented in Section 3. Finally, conclusions are provided in Section 4.

2. METHODOLOGY

2.1. Overview of sorted index numbers (SIN) method

The proposed method utilizes sorted index numbers other than the original facial features as templates for recognition. The procedure of creating the proposed SIN feature vector is as follows:

1. Extract feature vector $\mathbf{w} \in \mathbb{R}^n$ from the input face image $\mathbf{z}$.
2. Compute $\mathbf{u} = \mathbf{w} - \bar{\mathbf{w}}$, where $\bar{\mathbf{w}}$ is the mean feature vector calculated from the training data.
3. Sort the feature vector $\mathbf{u}$ in descending order, and store the corresponding index numbers in a new vector $\mathbf{g}$.
4. The generated vector $\mathbf{g} \in \mathbb{Z}^n$ that contains the sorted index numbers is stored as template for recognition.

For example, given $\mathbf{u} = \{u_1, u_2, u_3, u_4\}$, the sorted vector in descending order is $\hat{\mathbf{g}} = \{u_4, u_2, u_3, u_1\}$, then the template is $\mathbf{g} = \{4, 2, 3, 1\}$.

The method for computing the similarity between two SIN vectors is as follows:

1. Given two SIN feature vectors $\mathbf{g} \in \mathbb{Z}^n$ and $\mathbf{p} \in \mathbb{Z}^n$, where $\mathbf{g}$ denotes the template vector, and $\mathbf{p}$ denotes the probe vector. Start from the first element g_1 of $\mathbf{g}$.
2. Search for the corresponding element in $\mathbf{p}$, i.e., $p_i = g_1$. Record $d_1 = i - 1$, where i is the index number in $\mathbf{p}$.
3. Eliminate the obtained p_i in the previous step from $\mathbf{p}$, and obtain $\mathbf{p}^1 = \{p_1, p_2, ..., p_{i-1}, p_{i+1}, ..., p_n\}$.
4. Repeat step 2 and 3 on the following elements of $\mathbf{g}$ until g_{n-1}. Record $d_2, d_3, ..., d_{n-1}$.
5. The similarity measure of $\mathbf{g}$ and $\mathbf{p}$ is computed as $S(\mathbf{g}, \mathbf{p}) = \sum_{i=1}^{n-1} d_i$.

Illustration example:

1. For two SIN feature vectors $\mathbf{g} = \{4, 2, 3, 1\}$ and $\mathbf{p} = \{3, 2, 1, 4\}$, we first search the 1^{st} element $g_1 = 4$, and find that $p_4 = 4$. Therefore $d_1 = 4 - 1 = 3$. Eliminate p_4 from $\mathbf{p}$ and we form a new vector of $\mathbf{p}^1 = \{3, 2, 1\}$.
2. Search the 2^{nd} element $g_2 = 2$, and find that $p_2^1 = 2$. Therefore $d_2 = 2 - 1 = 1$. Eliminate p_2^1 from $\mathbf{p}^1$ and form a new vector of $\mathbf{p}^2 = \{3, 1\}$.
3. Search the 3^{rd} element $g_3 = 3$, and find that $p_1^2 = 3$. Therefore $d_3 = 1 - 1 = 0$.
4. Compute $S(\mathbf{g}, \mathbf{p}) = \sum_{i=1}^{n-1} d_i = 3 + 1 + 0 = 4$.

2.2. Methodology analysis

To understand the underlying rationale of the proposed algorithm, we first look into an alternative presentation of the method, named Pairwise Relational Discretization (PRD). The procedure of producing the PRD feature vector is as follows:

1. Extract feature vector $\mathbf{w} \in \mathbb{R}^n$ from the input face image $\mathbf{z}$.
2. Compute $\mathbf{u} = \mathbf{w} - \bar{\mathbf{w}}$, where $\bar{\mathbf{w}}$ is the mean feature vector calculated from the training data.
3. Compute binary representation of $\mathbf{u}$ by comparing the pairwise relation of all the elements in $\mathbf{u}$ according to:
$$b_{ij} = \begin{cases} 1 & u_i \geq u_j; \\ 0, & u_i < u_j; \end{cases}$$
4. Concatenate all the generated binary bits into one vector $\mathbf{b} = \{b_{12}, ..., b_{1n}, b_{23}, ..., b_{2n}, b_{34}, ..., b_{n-1,n}\}$. Store the binary vector $\mathbf{b}$ as template for recognition.

The similarity measure of the PRD method is Hamming distance. Unlike traditional discretization method, which quantizes individual elements based on some predefined quantization levels, the proposed method takes the global characteristics of the feature vectors into consideration. This is interpreted by comparing the pairwise relation of all groups of two elements in the vector. The intuition behind the idea is to take a n-dimensional space as superposition of all combinations of 2-dimensional planes. In n-dimensional subspace, when the similarity of two vectors is evaluated by Euclidean distance (i.e. spatial closeness), each element of the vectors are treated as coordinates in the corresponding basis. The elements are essentially the projection coefficients of the vector onto each basis (i.e. lines). Here, instead of projecting onto lines, we explore the projection onto 2-D planes. The discretization step partitions a plane into two regions by comparing the pairwise relation. It reduces the sensitivity of the variation of individual elements, and therefore possibly provides better error tolerance. The mean centralization step is to leverage the significance of each element such that no single dimension will overwhelm others. For two points in n-dimensional subspace, if they are spatially close to each other, then in large number of 2-D planes, their projection location should be close to each other, i.e., small Hamming distance, and vise versa.

A major drawback of the PRD method is the high dimensionality of the generated binary PRD vector. For a n-dimensional vector, the generated binary vector $\mathbf{b}$ will have a size of $\frac{n(n-1)}{2}$. This problem introduces high storage and computational requirements. To improve this, we note that the PRD method is based on pairwise relation of all the vector elements, and the same information can be exactly preserved by storing the sorted index numbers of the vector, i.e., any single bit in $\mathbf{b}$ can be derived from the stored SIN vector. Let $\mathbf{g}$ and $\mathbf{p}$ denote the SIN vectors of template and probe images respectively, then there are $n - 1$ bits that are associated with the first element g_1 in $\mathbf{b}_g$ (corresponding PRD vector of template image). Searching for the corresponding element in $\mathbf{p}$, i.e., $p_i = g_1$, then all the index numbers to the left of p_i will have different bit values in corresponding PRD vector $\mathbf{b}_p$, i.e., $d_1 = i - 1$ corresponds to the Hamming distance of the bits in $\mathbf{b}_g$ and $\mathbf{b}_p$ that are associated with index number g_1. It should be noted that since the Hamming distance for all the bits associated with $p_i = g_1$ have been computed, the p_i element should be removed for the calculation of next iteration. After the Hamming distances for all the elements in $\mathbf{g}$ and $\mathbf{p}$ are computed, the sum of them will correspond to the Hamming distance of $\mathbf{b}_g$ and $\mathbf{b}_p$. Let $H(\mathbf{b}_1, \mathbf{b}_2)$ denotes the Hamming distance between two binary vectors $\mathbf{b}_1$ and $\mathbf{b}_2$, then $H(\mathbf{b}_g, \mathbf{b}_p) = S(\mathbf{g}, \mathbf{p}) = \sum_{i=1}^{n-1} d_i$. Therefore, the proposed SIN and PRD methods produce exactly the same results. To test the effectiveness of SIN over PRD in computational complexity, we performed experiments on a computer with Intel $Core^{TM}2$ CPU 2.66GHz. With an original feature vector of dimensionality 100, the average time for PRD feature extraction and matching is 26.2 ms, while the SIN method only consumes less than 0.9 ms.

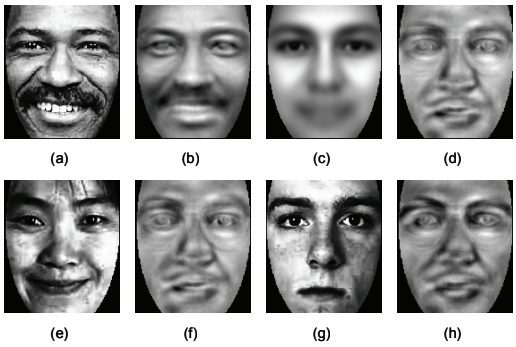

Fig. 1. Comparison of original and reconstructed images.

2.3. Privacy preserving property

Since the proposed method only stores the index numbers of the sorted feature vector $\mathbf{u}$, the transformation from $\mathbf{u}$ to the corresponding SIN vector $\mathbf{g}$ is non-invertible. There is no effective approximate reconstruction being possible to recover the values of $\mathbf{u}$ from $\mathbf{g}$. The most an adversary can do is to estimate the values of $\mathbf{u}$ based on some statistics or his/her own features. By using such method, an adversary can only produce a distorted version of the original image. To provide some insight into the privacy protection property of the proposed method, we compare the reconstructed image with the original image through different methods in Fig. 1.

Fig. 1-a shows an image $\mathbf{z}$ and Fig. 1-b is the reconstructed image from its first 100 PCA coefficients $\mathbf{u}$. The reconstruction is performed by $\widehat{\mathbf{z}} = \Psi(\mathbf{u} + \Psi^T\bar{\mathbf{z}})$, where Ψ is the PCA projection matrix, and $\bar{\mathbf{z}}$ is the mean image obtained from the training set. It is obvious that the PCA approach can not preserve privacy since the original visual information is very well approximated. Since the mean image is usually stored in the system database, an adversary may obtain and use it for the reconstruction of images. Let $\mathbf{g}$ be the stored SIN vector of image $\mathbf{z}$, the reconstruction can be performed by extracting PCA coefficients from the mean image $\bar{\mathbf{z}}$, sort and map to the corresponding element in $\mathbf{g}$, and then perform the same reconstruction procedure as above. Fig. 1-d shows the reconstructed image from the mean image in Fig. 1-c using the SIN vector $\mathbf{g}$ of image $\mathbf{z}$. It can be seen that the reconstructed image has a large distortion from the original one. Alternatively, an adversary can use the features of himself/herself to reconstruct $\mathbf{z}$. Fig. 1-f and 1-h are the reconstructed images from the PCA coefficients of images in Fig. 1-e and 1-g respectively, using the SIN vector $\mathbf{g}$ of image $\mathbf{z}$, which also demonstrate large distortion from the original image. An interesting observation is that although using different PCA coefficients, the reconstructed images in Fig. 1-d, 1-f, and 1-h are quite similar. This further demonstrates that other than the detailed feature values, the sorted index numbers indeed contain discriminant information.

The above analysis shows that the privacy of the users can be protected by using the sorted index numbers. This is demonstrated through the visual dissimilarity of the original and reconstructed images. The binary nature of the PRD vector makes it a candidate for cryptographic key generation in biometric crypto-systems. This can be carried out by using the fuzzy commitment scheme [2] alike techniques, where the binary vector is bounded with a randomly generated key, and error correction algorithms are applied for error tolerance. On the other hand, the SIN method can also be applied in conjunction with intentional transformation methods for changeable biometrics generation. This can be achieved by applying distance preserving transformations (such as random orthogonal transformation) prior to the sorting operation. In this paper, we will focus on demonstrating the effectiveness of the SIN/PRD method on original features only.

3. EXPERIMENTAL RESULTS

To approach more realistic face recognition applications, this paper evaluate the performance of the proposed method on a generic data set, in which the intrinsic properties of the human subjects are trained from subjects other than those to be recognized. The generic data set was initially organized for the purpose of demonstrating the effectiveness of the generic learning framework [11]. It contains 5676 images of 1020 subjects from 5 well-known databases, FERET [12][13], PIE [14], AR [15], Aging [16], and BioID [17]. All images are aligned and normalized based on the coordinate information of some facial feature points. The details of image selection and configuration can be found in [11]. To study the effects of different feature extractors on the performance of the proposed method, we compare Principal Component Analysis (PCA) and Kernel Direct Discriminant Analysis (KDDA) [18], which have been demonstrated to be effective appearance based approaches for face recognition [11].

3.1. Experimental results on face identification

For face identification, we use all the 5676 images in the data set for experiments. A set of 2836 images from 520 human subjects were randomly selected for training, and the rest of 2840 images from 500 subjects for testing. There is no overlap between the training and testing subjects and images. The test is performed on an exhaustive basis, such that each time, one image is taken from the test set as probe image, while the rest of the images in the test set as gallery images. This is repeated until all the images in the test set were used as the probe once. The classification is based on nearest neighbor.

Table 1 compares the correct recognition rate (CRR) of SIN method with Euclidean and Cosine distance measures at different feature dimensions. It can be observed that at higher dimensionality, the SIN method may boost the recognition accuracy of PCA significantly, while maintain the good performance of the stronger feature extractor KDDA. The PCA method projects images to directions with highest variance, but not the discriminant ones. This will become more severe in large image variations due to illumination, expression, pose and aging. When computing the similarity between two PCA vectors, the distance measure is sensitive to the variation of individual element, particularly those directions corresponding to noise. The SIN method, on the other hand, reduces this sensitivity by comparing the relative relation of the projections, and therefore possibly provides better error tolerance. In the case of strong extractors such as KDDA, the SIN method will approximate the distance between two vectors, and hence preserves the recognition accuracy.

3.2. Experimental results on face verification

For face verification, we exclude image samples with large pose variation ($> 15^\circ$), and select 4666 images from 1020 subjects as the

Dim.	PCA			KDDA		
	Euc.	Cos.	SIN	Euc.	Cos.	SIN
20	56.30	56.31	52.32	40.04	41.09	34.86
40	60.09	61.09	61.94	61.44	65.28	61.94
60	63.52	62.96	66.06	71.73	74.86	74.68
80	64.37	64.44	68.84	81.76	83.27	81.76
100	65.14	65.18	71.27	79.05	80.42	80.07

Table 1. Face identification results (in %).

Dim.	PCA			KDDA		
	Euc.	Cos.	SIN	Euc.	Cos.	SIN
20	20.05	19.23	13.78	25.22	20.42	20.97
40	19.09	17.81	11.46	21.49	16.22	14.54
60	18.52	17.42	10.28	18.80	13.41	10.97
80	18.50	17.15	9.72	10.96	9.90	7.19
100	18.20	16.94	9.46	10.41	8.84	6.52

Table 2. Obtained EER (in %) for face verification.

verification data set. In our experiments, we randomly select 2388 images from 520 subjects as the training set, and 2278 images of the rest 500 subjects as the testing set. There is no overlap between the training and the testing subjects and images. The evaluation was also performed on an exhaustive basis, where every single image is used as a template once, and the rest of the images in the test set as the probe images.

Table 2 compares the obtained EER (equal error rate, operating point where false accept rate and false reject rate are equal) of SIN with Euclidean and Cosine distance at different dimensions when PCA and KDDA are used as feature extractors. In general, the Cosine metric outperforms the Euclidean distance measure, and the proposed SIN method improves both the verification accuracy of PCA and KDDA at almost all dimensions. This further demonstrates that the sorted index numbers indeed offer better error tolerance and provide more discriminant representation.

4. CONCLUSION

This paper introduced a novel approach for addressing the challenging problem of privacy preserving face recognition. The proposed method stores the sorted index numbers of facial feature vectors as biometric template for recognition. A similarity measure is introduced for computing the distance between two SIN vectors. Experimental results on a large and complex data set demonstrate that the proposed solution may improve the recognition accuracy in both identification and verification scenarios. It is shown that the transformation from original features to the sorted index numbers is non-invertible. Since there is no effective method of reconstructing the original features as well as images, the proposed method provides privacy protection. The proposed approach may also be combined with cryptographic techniques and intentional transformation based methods to produce changeable biometric template. Although we focus on face recognition problem in this paper, the proposed method is general, and it is expected that such method can also be applied to other biometrics.

5. REFERENCES

[1] U. Uludag, S. Pankanti, S. Prabhakar, and A. K. Jain, "Biometric Cryptosystems: Issues and Challenges", Proc. of the IEEE, vol. 92, no. 6, pp. 948-960, 2004

[2] A. Juels, and M. Wattenberg, "A fuzzy commitment scheme", Proc. of sixth ACM Conf. on Computer and Communication Security, pp. 28-36, 1999

[3] A. Juels, and M. Sudan, "A fuzzy vault scheme", Proc. of IEEE International Symp. on Information Theory, pp. 408, 2002

[4] Y. Dodis, L. Reyzin, A. Smith, "Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data", Proc. Eurocrypt 2004, pp. 523-540, 2004.

[5] T. A. M. Kevenaar, G. G. Schrijen,M. van der Veen, A. H. M. Akkermans, and F. Zuo, "Face recognition with renewable and privacy preserving binary templates", Fourth IEEE Workshop on Auto. Identi. Adv. Tech. Oct. 2005 Page(s):21 - 26

[6] R. M. Bolle, J. H. Connel, N. K. Ratha, "Biometric perils and patches", Pattern Recognition, vol. 35, pp. 2727-2738, 2002

[7] A.B.J. Teoh, D.C.L Ngo and A. Goh, "BioHashing: two factor authentication featuring fingerprint data and tokenised random number", Pattern Recognition, vol. 37, pp. 2245-2255, 2004.

[8] M. Savvides, B. V.K. Vijaya Kumar, and P. K. Khosla, "Cancelable biometric filters for face recognition", Proc. of the 17th Int. COnf. on Pattern Recognition, pp. 922-925, 2004

[9] W. J. Scheirer, T. E. Boult, "Cracking Fuzzy Vaults and biometric encryption", Biometrics Symposium, BSYM 2007, Baltimore, Maryland, USA, September, 2007.

[10] R. Brunelli, T. Poggio, "Face recognition: feature versus templates", IEEE Trans. Pattern Anal. Mach. Intell. 15(10) 1042-1052, 1993

[11] J. Wang, K. N. Plataniotis, J. Lu and A. N. Venetsanopoulos, "On Solving the Face Recognition Problem with One Training Sample per Subject", Pattern recog. 39(2006), pp. 1746-1762

[12] P. J. Phillips, H. Wechsler, J. Huang, P. Rauss, "The FERET database and evaluation procedure for face recognition algorithms", Image Vision Comput. J. 16(5), 1998, pp. 295-306

[13] P. J. Phillips, H. Moon, P. J. Rauss, and S. Rizvi, "The FERET evaluation methodology for face recognition algorithms", IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 22, No. 10, October 2000.

[14] T. Sim, S. Baker, and M. Bsat, "The CMU Pose, Illumination, and Expression Database", IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 25, No. 12, December, 2003, pp. 1615 - 1618.

[15] A. M. Martinez, R. Benavente, The AR face database, CVC Technical report 24, 1998

[16] Aging Database, (http://sting.cycollege.ac.cy/ alanitis/fgnetaging/).

[17] BioID Database, (http://www.humanscan.de/support/downloads/facedb.php).

[18] Juwei Lu, K.N. Plataniotis, and A.N. Venetsanopoulos, "Face Recognition Using Kernel Direct Discriminant Analysis Algorithms", IEEE Trans. on Neural Networks, Vol. 14, No. 1, Page: 117-126, January 2003.