PHYSICS-BASED REVOCABLE FACE RECOGNITION

Gaurav Aggarwal[†]*, Nalini K. Ratha*[††]*, Jonathan H. Connell*[††] *and Ruud M. Bolle*[††]

[†]Center for Automation Research
University of Maryland, College Park
gaurav@cfar.umd.edu

[††]IBM T. J. Watson Research Center
Hawthorne, NY, USA
{ratha, jconnell, bolle}@us.ibm.com

ABSTRACT

We present a face reconstruction approach for revocable face matching. The proposed approach generates photometrically valid cancelable face images by following the image formation process. Given a face image, the approach estimates facial albedo followed by a subject-specific key based photometric deformation to generate a cancelable face image. The proposed approach allows for using any available face matcher to perform verification or recognition in the transformed domain, a capability missing from most exisiting works on cancelable face matching. Experiments are performed to evaluate the performance, privacy and cancelable aspects of the face images reconstructed using the approach. Results obtained are very promising and make a strong case for such backward compatible cancelable face representations that can seamlessly make use of advancements in automatic face recognition research.

Index Terms— cancelable biometrics, face recognition, signal-level transforms

1. INTRODUCTION

Biometrics refers to the measurement of physical or behavioral traits of humans, often targeted towards the goal of verifying or determining personal identity [1]. Biometric characteristics provide a better alternative to establish identity of a person as compared to PINs, passwords, etc., which can easily be stolen or passed on to others fairly easily. A lot of research has been done to recognize humans based on their fingerprints, faces, gait, iris, etc. Advancements of research in this area has made it possible to visualize biometric systems that can operate in real world scenarios.

The advancement and popularity of biometric systems has brought concerns of *biometric-theft*. Unlike PINs or passwords, which can be changed at will when compromised, biometric traits are unique and permanent. This leads to the observation that though biometrics are authentic, they are not secure (or private like passwords). If compromised, biometric signatures cannot be revoked or canceled. It allows for rogue establishments to track subjects across databases and institutions without consent.

The concern of biometric privacy has led to research efforts to secure biometrics [2]. One popular way is to combine biometrics with user-provided keys or passwords to make them secure. The user-specific private key is used to encrypt biometric template which is stored in the database. The encrypted template stored in the database is used for further matching. For matching purposes, the same encryption scheme is used to transform the query template to compare it with the stored secure template. Quite clearly, such an approach combines the advantages of biometric based authentication and password-based privacy and revocability.

Ratha *et al.* [2], in their pioneering work, present several one way (non-invertible) transforms for constructing multiple secure identities from a fingerprint. They show that a user can be given as many biometric identifiers as needed by issuing a new transformation key which can be canceled and replaced when compromised. Savvides *et al.* [3] extend their earlier work on correlation filter based face matching to produce cancelable biometric representations. They show that convolving the training images with any random convolution kernel before building the filter does not change the resulting correlation output peak-to-sidelobe ratios, thus preserving the authentication performance while maintaining privacy. Boult [4] introduces robust biometric transform that can be used for revocable face authentication. The transformed feature vector is separated into a fractional and an integral part where the integral part is encrypted while the other is left unsecured. Teoh *et al.* [5] elaborate biometric-hash framework by integrating biometric and user-specific password using Random Multispace Quantization (RMQ). The process is carried out by first obtaining a fixed length feature vector form the input biometric followed by a non-invertible random subspace projection and quantization.

One of the main problems in encryption-based biometric authentication approaches is that they tend to be sensitive to variability/noise in the input biometric space. Inherently, biometrics show a great deal of intra-class variability either due to natural causes or external imaging conditions. It is difficult to design an encryption scheme that can suitably transform features extracted from such input data minimizing within-class scatter as compared to the between-class scatter. Unlike input biometric space, in which one can perform some sort of learning to account for such intra-class variabilities, such learning is not easy in the encrypted space. Another drawback of encrypting feature extracted from the input biometrics is that such approaches tend to be specific to the features used. Therefore, it may not always be easy for such approaches to take advantage of the new developments in the field of biometric matching.

In this paper, we propose a physics-based face reconstruction approach that addresses these issues for cancelable face matching. Given an input face image, the proposed technique reconstructs a transformed face image that can be matched using any publicly available matcher. Depending on the capability of the face matcher used to compare the reconstructed face images, the variability/noise in the input biometric can be accounted for even though matching is performed in the transformed domain.

1.1. Organization of the paper

The paper is organized as follows. The proposed face reconstruction approach is described in Section 2. Results of extensive experimental evaluations performed to validate the usefulness of the approach in terms of privacy, security and matching performance are shown in

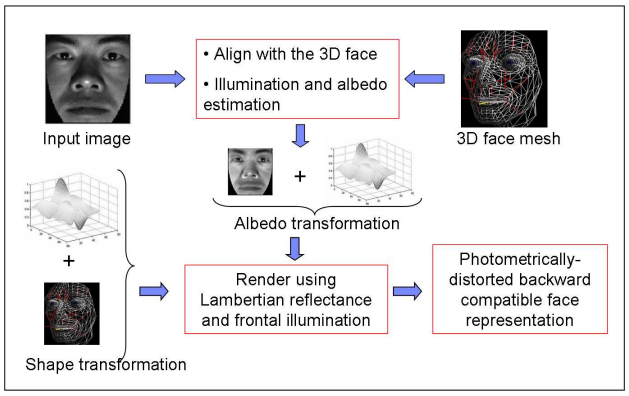

Fig. 1. A schematic of the proposed approach.

Section 3. The paper concludes with a brief summary and concluding remarks in Section 4.

2. PHYSICS-BASED FACE RECONSTRUCTION

An input face image is the result of an interplay between the physical characteristics of a real 3D face, external imaging environment (illumination, view, etc.), capturing device, etc. Our goal is to create another face image from the input image that can be used as a cancelable representation of the face, that can be matched using any available face matcher. One of the critical components of such an approach is appropriateness of the transformation from the input face to the desired cancelable face image such that the output is photometrically valid. Quite clearly, direct manipulation of the image intensity values may lead to images which are physically unrealizable.

In this paper, we first estimate albedo from a single input face image. This is followed by user-specific key based transformation of albedo. Due to the absence of real 3D shape information of the input face and the difficulty in estimating shape from a single image, we use a transformed (distorted) version of the average facial 3D shape. As with albedo, the kind and amount of shape distortion is guided by the user-specific key. Once we have the distorted albedo and shape, we render a face image that does not reveal the identity of the subject in the input image. The vast range of possible transformations (or distortions) on estimated albedo and 3D facial shape provides cancelability to the approach for scenarios when the template is compromised. Fig. 1 shows a schematic of the proposed approach. The various steps of the proposed approach are described as follows.

2.1. Albedo estimation

The first step of our approach is to estimate surface albedo from the input face image. Without loss of generality, face images are assumed to be pre-cropped and pose-normalized to be in the frontal pose. Albedo estimation is performed using the non-stationary stochastic filtering framework proposed by Biswas *et al.* [6]. Given a coarse albedo map (obtained using the average facial 3D shape of humans), the approach estimates a more robust albedo map by accounting for the statistics of errors in surface normal and light source estimation in an image restoration framework [7]. Readers are encouraged to read [6] for technical details.

2.2. Albedo and shape transformation

In this step, the estimated albedo and the average average facial 3D shape is transformed using the user-provided secure key. There can be a large number of choices for such a transformation, we use one based on mixture of Gaussians. Albedo is transformed by multiplying it with a mixture of Gaussian image. The number, peak locations, and variance of the Gaussian distributions is determined using the key. For shape transformation, we generate another mixture of Gaussians surface and linearly combine it with the average 3D facial shape. As with albedo, the user-specific key determines the specifics of the mixture of Gaussians surface.

2.3. Image Reconstruction

The transformed albedo and shape are used to reconstruct a photometrically valid face image. Assuming Lambertian reflectance model, the desired image can easily be generated using the following relation

$$\mathbf{I}_r = \rho_r \max(\mathbf{n}_r \cdot \mathbf{s}, 0) \tag{1}$$

where $\mathbf{I}_r$ is the reconstructed transformed face image, ρ_r is the transformed albedo map, $\mathbf{n}_r$ is the transformed surface normal map and $\mathbf{s}$ is the light source direction which is taken to be $[0, 0, 1]^T$ for frontal lighting. Fig. 2 shows a few images generated using this approach.

3. EXPERIMENTAL EVALUATION

In this section, we describe the experiments performed to evaluate the usefulness of the proposed backward-compatible cancelable face reconstruction. In our implementation, the user-defined keys are generated using a random number generator that defines the number (5-10), location and variance of Gaussian peaks required to generate distorted images. The experiments are performed on illumination part of the PIE face dataset [8] that consists of face images of 68 subjects under 21 challenging illumination conditions (Fig. 3). Each experiment consists of matching images in one illumination scenario against another. This results in 68 genuine and 68×67 impostor pairs. All the verification results and score distributions presented in this paper are obtained by repeating the experiment for all $\binom{21}{2}$ pairs of illumination conditions, thereby resulting in $\binom{21}{2} \times 68$ genuine and $\binom{68}{2} \times 68 \times 67$ impostor pairs. In addition to evaluating privacy and revocability, experiments also reflect the illumination-invariance property of the approach. Illumination-invariance is a byproduct of using albedo images as opposed to direct intensity images for transformed face reconstruction. In all the experiments, similarity scores are computed using Principal Component Analysis (PCA). The PCA bases are learnt from FRGC training data [9] that consists of 366 training face images.

3.1. Performance

Fig. 4 shows the genuine and impostor score distributions obtained using the reconstructed faces. In this experiment, every subject has a different transformation key. The plot shows the distributions obtained in two different runs of the experiment using different sets of keys for each identity. The genuine/impostor distributions hardly overlap leading to almost flawless performance. Note that the proposed approach is able to account for illumination variations present in the original images, a capability missing in most previous cancelable face matching approaches.

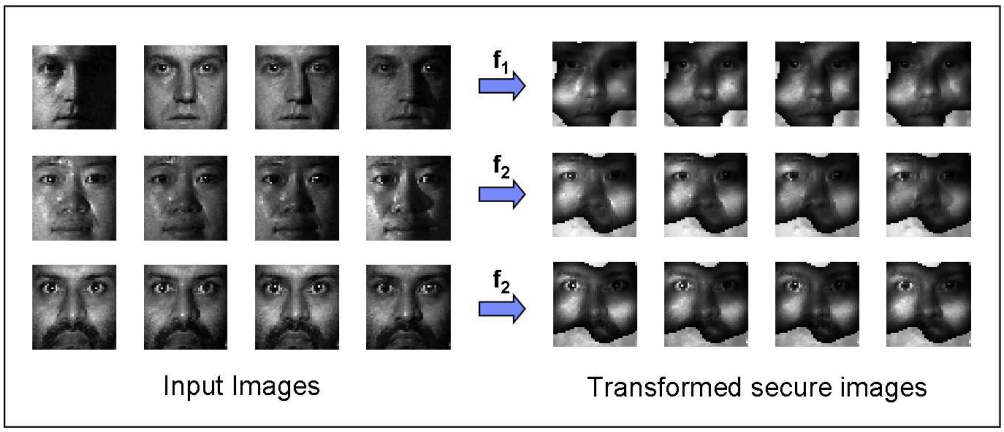

Fig. 2. Examples of transforms applied to a few images from the PIE dataset.

Fig. 3. The 21 illumination conditions in the PIE dataset.

3.2. Lost key scenario

We now evaluate the performance of the approach by using same transformation key for all the subjects. This simulates the stolen/lost key scenario when an adversary somehow gets hold of a users key and tries to break into the system using that key. Fig. 5 shows that the separation between the genuine/impostor distributions is preserved even when same transformation key is used for all the subjects. In fact, the reconstructed faces perform better in a verification setting as compared to original input images even when same transformation key is used for reconstruction (Fig. 6).

3.3. Privacy and revocability

We first compare the genuine and impostor score distributions obtained while matching reconstructed images against the original input images , i.e., using original images in the gallery while the reconstructed images as queries. The experiment is repeated by replacing original images with another set of transformed images generated using a different set of keys. Fig. 7 shows that the genuine and impostor score distributions have hardly any separation indicating that the reconstructed faces reveal hardly any identifying information when compared against the original or (differently) transformed images. To further evaluate privacy/revocability of the proposed approach, we also compare the mated score distributions obtained while matching 1) original images against transformed images (should be low for privacy), 2) transformed images against other transformed images generated using the same key (should be high for good performance),

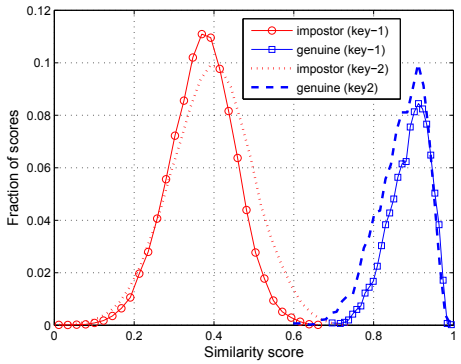

Fig. 4. Impostor and genuine score distributions obtained using the generated face images. The plot shows results obtained in two different runs of the proposed algorithm using different set of keys.

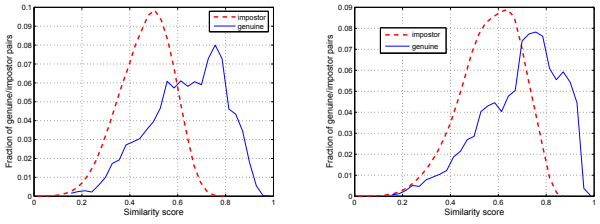

Fig. 5. Lost key scenario: Genuine/impostor score distributions obtained in matching experiments on the face images reconstructed using the same key for all identities (left) and the original input images (right). The genuine/impostor separation is preserved even when same key is used to transform all identities.

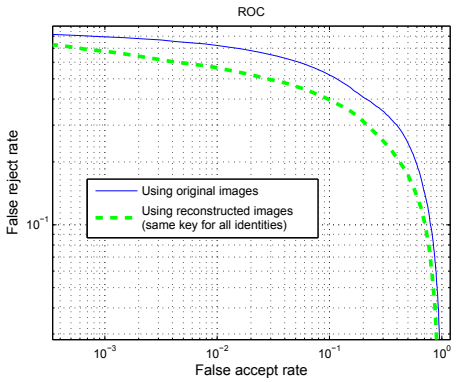

Fig. 6. Lost key scenario: Comparison of Receiver Operator Characteristic (ROC) curves obtained in a verification experiment with the original images in the gallery while the transformed faces (generated using same key for all identities) as queries.

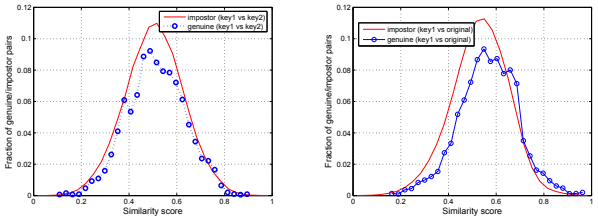

Fig. 7. Privacy/revocability test: 1) Genuine/impostor score distributions obtained using transformed image set 1 in the gallery while transformed image set 2 as queries (left), and 2) Genuine/impostor score distributions obtained using the original images in the gallery while the transformed ones as the queries (right).

and 3) transformed images against other transformed images generated using a different key (should be low for revocability). Fig. 8 shows that the genuine score distributions are in fact as desired proving the privacy and revocability aspects of the proposed approach.

4. SUMMARY

Unlike inter-operable fingerprint templates, there is no common format for face features other than the image itself. In order to achieve backward compatibility, we proposed a physics-based face reconstruction approach for cancelable face matching. Given an input face image, the proposed technique reconstructs a new transformed face image that can be matched using any available matcher. We tested our approach using a standard database with several different transforms. The results are extremely encouraging. We will test the scalability of our approach using larger databases and publicly available face matchers in future. Note that though it is impossible for an adversary to get the original image back from just a transformed image and the corresponding key, the transform is invertible if he/she has access to the exact distortion algorithm used to obtain the transformed images.

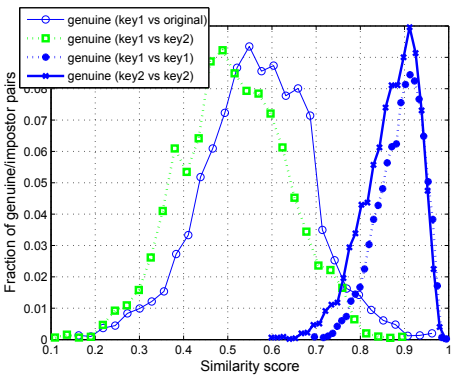

Fig. 8. Privacy/revocability test. Comparison of distributions of mated scores: 1) Original image against transformed image (should be low for privacy), 2) Transformed image against other transformed image generated using the same key (should be high for good performance), 3) Transformed image against other transformed image generated using a different key (should be low for revocability).

5. REFERENCES

[1] R. Bolle, J. H. Connell, S. Pankati, N. K. Ratha, and A. W. Senior, *Guide to Biometrics*, Springer verlag, 2003.

[2] N. K. Ratha, S. Chikkerur, J. H. Connell, and R. M. Bolle, "Generating cancelable fingerprint templates," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 4, pp. 561–572, 2007.

[3] M. Savvides, B. V. K. Vijayakumar, and P. K. Khosla, "Cancelable biometrics filters for face recognition," in *Proceedings of International Conference on Pattern Recognition*, 2004, vol. 3, pp. 922–925.

[4] T. Boult, "Robust distance measures for face-recognition supporting revocable biometric tokens," in *Proceedings of International Conference on Automatic Face and Gesture Recognition*, 2006, vol. 3, pp. 560–566.

[5] A. B. J. Teoh, A. Goh, and D. C. L. Ngo, "Random multispace quantization as an analytic mechanism for biohashing of biometric and random identity inputs," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 12, pp. 1892–1901, 2006.

[6] S. Biswas, G. Aggarwal, and R. Chellappa, "Robust estimation of albedo for illumination-invariant matching and shape recovery," in *Proceedings of IEEE International Conference on Computer Vision*, October 2007, pp. 1–8.

[7] H. C. Andrews and B. R. Hunt, *Digital Image Restoration*, Prentice-Hall signal processing series, 1977.

[8] 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, pp. 1615–1618, 2003.

[9] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek, "Overview of the face recognition grand challenge," in *Proceedings of International Conference on Computer Vision and Pattern Recognition*, 2005.