A Cancellable and Privacy-Preserving Facial Biometric Authentication Scheme

Tyler Phillips[*], Xukai Zou[*], and Feng Li[**]

[*]Department of Computer and Information Science, Indiana University-Purdue University Indianapolis
[**]Department of Computer and Information Technology, Indiana University-Purdue University Indianapolis

***Abstract*—In recent years, biometric, or "who you are," authentication has grown rapidly in acceptance and use. Biometric authentication offers users the convenience of not having to carry a password, PIN, smartcard, etc. Instead, users will use their inherent biometric traits for authentication and, as a result, risk their biometric information being stolen. The security of users' biometric information is of critical importance within a biometric authentication scheme as compromised data can reveal sensitive information: race, gender, illness, etc. A cancellable biometric scheme, the "BioCapsule" scheme, proposed by researchers from Indiana University Purdue University Indianapolis, aims to mask users' biometric information and preserve users' privacy. The BioCapsule scheme can be easily embedded into existing biometric authentication systems, and it has been shown to preserve user-privacy, be resistant to several types of attacks, and have minimal effects on biometric authentication system accuracy.**

In this research we present a facial authentication system which employs several cutting-edge techniques. We tested our proposed system on several face databases, both with and without the BioCapsule scheme being embedded into our system. By comparing our results, we quantify the effects the BioCapsule scheme, and its security benefits, have on the accuracy of our facial authentication system.

Index Terms*—biometric authentication, cancellable biometrics, facial authentication, BioCapsule scheme

I. Introduction

A. Biometric Authentication

There are three types of authentication: "what you know," "what you have," and "who you are."

"What you know" authentication is the most popular type of authentication, and it deals with a user possessing secret information that they will provide to the system in order to be authenticated. Passwords, pass-phrases, PIN numbers, and patterns are all examples of "what you know" authentication.

"What you have" authentication is another commonly used type of authentication where users must possess an object to provide to the system in order to be authenticated. Smart cards and several software dongles, such as the iLok, are examples of types of "what you know" authentication.

The final type of authentication is "who you are" authentication. This type of authentication deals with characteristics, physiological or behavioral, intrinsically linked with a user. Some examples of such characteristics are: fingerprint, iris, face, ECG, gait, speech, and keystroke patterns. Information used to quantify these characteristics, or biometrics, is provided to the system in order to be authenticated.

"Who you are" authentication is quickly growing in acceptance and use in a wide variety of domains. One major advantage of "who you are" authentication is the convenience it offers to users. Unlike the other two types of authentication, "who you are" authentication does not require users to carry their means of authentication. Rather than provide some carried knowledge or physical object, users provide the biometrics intrinsically linked to them. Their biometric information is sampled, transformed into a biometric template, and compared to registered biometric templates of who the user claims to be in order for an authentication decision to be made.

Though this "who you are" authentication method offers the user convenience, it also poses privacy and security threats to the user. If a user's credentials are stolen within a "what you have" or "what you know" authentication scheme, the stolen credentials can be revoked and replaced. Examples of this are a password being stolen and then reset or a credit card being stolen, cancelled, and replaced. When intrinsically linked biometrics are used as credentials and are stolen, there is no reasonable way to revoke and reset the biometric credentials. This would require a user to make major changes to their physiological or behavioral traits. In addition to not being revocable, stolen biometric credentials can reveal sensitive information about users such as: race, gender, illness, etc.

Several cancellable biometric schemes have been proposed to allow users to revoke and recreate stolen biometric templates. They typically work by performing some alterations to a user's biometric template. The altered template is compared with other altered templates to be authenticated. This way, if the user's altered biometric template is stolen, the altered template can be revoked by the user. The user can then alter their biometric template in a different way and continue using the system.

B. BioCapsule Scheme

One promising cancellable biometric scheme, the "BioCapsule" (BC) scheme, has been proposed by researchers [1]. The BC scheme involves fusing user biometrics with the biometrics of a reference subject (RS) in order to preserve user privacy and offer cancellability of stolen biometric templates. Since the BC method works by fusing user and RS biometrics, it can easily be embedded into existing biometric authentication systems where preprocessing and feature extraction schemes may vary.

This is the author's manuscript of the article published in final edited form as:
Phillips, T., Zou, X., & Li, F. (2017). A Cancellable and Privacy-Preserving Facial Biometric Authentication Scheme. In 2017 IEEE 14th International Conference on Mobile Ad Hoc and Sensor Systems (MASS) (pp. 545–549). https://doi.org/10.1109/MASS.2017.80

TABLE I
FACE DATABASE ATTRIBUTES

Database	Number of Subjects	Images per Subject	Image Size	Var. in Expression	Var. in Illumination	Var. in Pose	Partial Face Images
JAFFE	10	20	256x256	Yes	No	No	No
AT&T	40	10	92x112	Yes	No	Yes	No
Yale	15	10	320x243	Yes	Yes	No	No
Faces95	72	20	180x200	Yes	Yes	Yes	Yes
BioID	19	19	896x592	Yes	Yes	Yes	Yes
IUPUI	18	10	200x260	Yes	Yes	Yes	Yes

After user biometrics are preprocessed and features are extracted, RS biometrics are also preprocessed and RS features are extracted. Once both a user and RS feature have been extracted, a signature is extracted from both features. Signature extraction works by moving a sliding window average across a feature matrix. After extracting a user and RS signature, we generate keys using both the signatures. The data within a signature is fed into a random number generator as seeds to generate values [0,1]. The resulting [0,1] random values are rounded to 0 and 1. Then we convert all 0 values to -1. This results in a key matrix of 1 and -1 values. After generating a user key from the user signature and generating a RS key from the RS signature, we are ready to perform fusion. The user feature is fused with the RS key through the use of the cross product. Likewise, the RS feature is fused with the user key through the cross product. Finally, the resulting two matrices are added together to generate a BC [1].

BCs will be registered to the system and associated with users and, at authentication time, user biometrics will be converted into a BC in order for authentication.

Researchers [1] have shown the BC scheme can offer users several security benefits such as:

- If a BC is compromised - Deriving the user or RS biometrics from the BC is equivalent to solving an undetermined equation.
- If many BCs are compromised - Deriving the RS used by a system through the use of many BCs is equivalent to solving an undetermined system of equations.
- If a BC and a RS are compromised - Security of user biometrics can be measured by the strength of the key used for fusion.
- If many users' biometrics and corresponding BCs are compromised - Deriving the RS used by the system through the use of many users' biometrics and corresponding BCs is NP-hard (very costly).
- If an attacker gathers a group of users' biometrics, the same users' BCs, and the RS biometrics - If the attacker has access to these tools and obtains another user's BC, using the RS biometrics and other data to find the new user's biometrics is NP-hard.

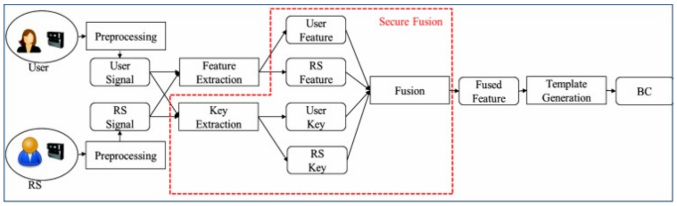

Fig. 1. BioCapsule Generation [1]

In addition to these security benefits, the BC method is privacy preserving and will mask the sensitive information associated with user biometrics (e.g. race, gender, illness, etc.).

Researchers [1] have also found that when the BC scheme is embedded into an existing iris biometric authentication scheme, the BC has little effect on the accuracy of the underlying system.

In this paper we propose a facial biometric authentication system which employs several cutting edge techniques. We will test our system with and without the BC scheme embedded to quantify the BC scheme's effect on system accuracy. The paper will be ordered as follows: in section two we discuss our datasets and system's preprocessing, feature extraction, and BioCapsule generation techniques; in section three we discuss our Euclidean distance experiment; in section four we offer concluding remarks; and in section five we outline future work plans and goals.

II. DATASETS, PREPROCESSING, FEATURE EXTRACTION, AND BIOCAPSULE GENERATION

A. Datasets

For our experiment, we perform a Euclidean distance test on six face databases: the JAFFE Database, the AT&T Face Database, the Yale Face Databse, the Faces95 Database, the BioID Face Database, and the IUPUI Face Database. Each of these databases, besides the IUPUI Face Database, are openly available online. The IUPUI Face Database was gathered using an implementation of our proposed system on a laptop. The databases range from very constrained to very unconstrained. We examined the attributes of each face database such as: source image size, variation in expression, variation in illumination, variation in pose, and partial face images (see Table 1). These attributes are known to pose challenges to face based biometric authentication systems as they complicate comparisons and often times lower system accuracy. Based on these attributes we rank the face database from most to least constrained in the following way:

1) JAFFE Face Database
2) AT&T Face Database
3) Yale Face Database
4) Faces95 Database

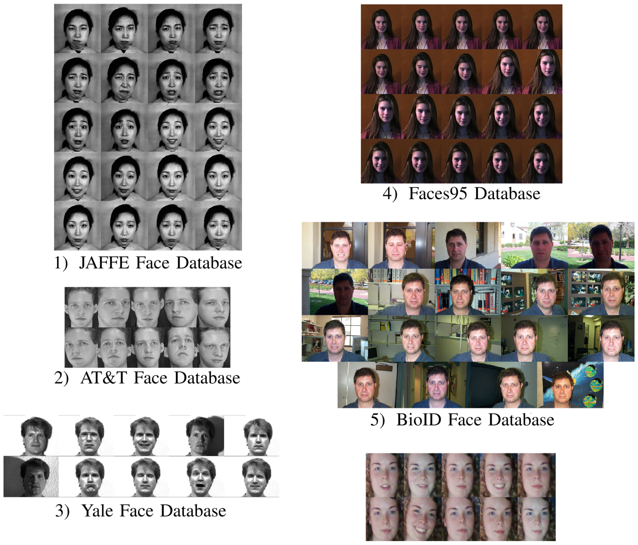

1) JAFFE Face Database

2) AT&T Face Database

3) Yale Face Database

4) Faces95 Database

5) BioID Face Database

6) IUPUI Face Database

Fig. 2. Sample Subject from each Database

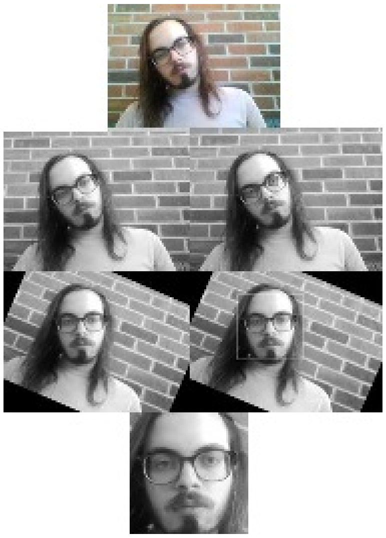

Fig. 3. Image Converted to Grayscale. Eyes are Detected then the Image is Rotated. Once Rotated, the Face is Detected and Cropped from the Image

5) BioID Face Database
6) IUPUI Face Database

We perform our experiment on all six databases to gauge how our authentication scheme works across different environments, with and without the BC scheme embedded into it.

B. Preprocessing and Normalization

Our authentication scheme begins with a number of preprocessing and normalization steps to prepare an image for feature extraction.

The first preprocessing step is to resize any input image to 196x180 pixels. This will normalize the size of any image input into the system.

Next, we reduce the image to a single color channel by simply converting the image to its grayscale representation.

Our third preprocessing step is to crop out the area of the grayscale image which contains a face. We crop the facial region of the image out of the surrounding image so that, once we are making authentication decisions, only faces are being compared rather than their surrounding environment. A simple tilt of a user's head can cause accuracy errors within a system when it is compared to images where the user's head is not tilted. To address this issue, we first try to detect the user's eyes within the grayscale image using Haar Cascade Classifiers. If two or more eyes can be detected within the image, we select the two most confident detections. We then find the center points of the two eye detections and the angle between them. We then rotate the image by this angle in order to make the resulting angle between the two eye detections 0°. After rotating the image, we attempt to perform facial detection using a second Haar Cascade Classifier. If we are able to detect a face within the rotated image, we crop out the detected facial region. If we cannot detect two eyes or a face within the rotated image, we attempt to detect a face within the original grayscale image. If a face can be detected, we crop the detected facial region from the image.

After we have cropped the detected facial region from the grayscale image, we perform our final preprocessing step, Tan-Triggs illumination normalization. Illumination poses a serious challenge in facial authentication systems. A slight variation in illumination between two similar images will cause differences in values of almost every pixel. To address this challenge we employ the Tan-Triggs illumination normalization technique as presented in [2]. Tan-Triggs works through gamma correction (enhancing differences within dark regions of an image while compressing the differences within the image's bright regions), Gaussian filtering, and contrast normalization. Researchers [2] have shown that the Tan-Triggs conversion increases facial recognition accuracy in challenging illumination conditions.

C. Feature Extraction

Once we have preprocessed and normalized an image, we are ready to perform feature extraction. Feature extraction involves retrieving representative data from the cropped, preprocessed face image. The retrieved data, or feature, is representative of the image in such a way that: we will be

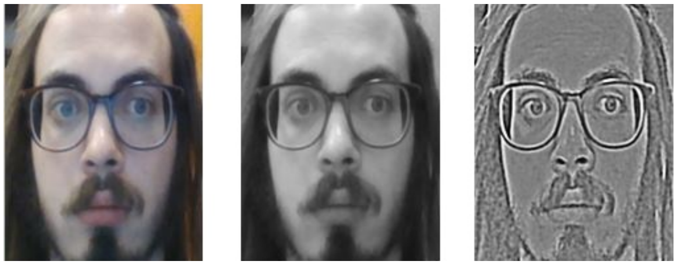

Fig. 4. Image Converted to Grayscale and Normalized using Tan-Triggs

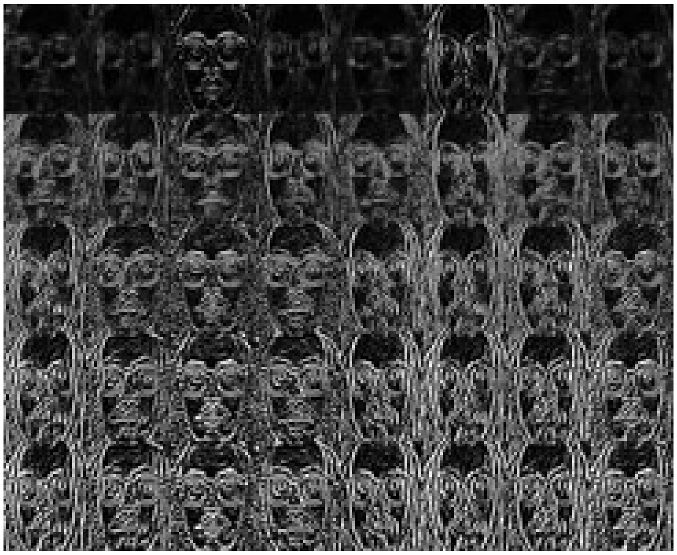

Fig. 5. Gabor Filterbank based Feature Extraction, based on Method of [3]

able to distinguish if the feature is being compared to a feature extracted from the face of the same person, or if it is being compared with a feature extracted from the face of a different person.

In order to perform feature extraction we apply a Gabor filterbank to the preprocessed, face image. A Gabor filter is a Guassian kernel function modulated by a sinosudal plane wave. Gabor filtering is a common texture analysis technique and is popular among feature extraction methods. We perform our Gabor filtering according to the parameters set within [3]. This method includes the use of a 40 filter Gabor filterbank of 8 orientations and 5 scales. We apply each of the 40 Gabor filters to a 196x180 preprocessed, face image. We then resize each of the resulting 196x180 Gabor filtered images to 32x30 pixels. We combine all 40 of the 32x30 Gabor filtered images into a single matrix, then reshape the matrix to a 38400x1 pixel feature vector.

D. BioCapsule Generation

To generate BioCapsules, we perform each of the steps of the BC scheme outlined in section one, using an extracted feature vector as input. Each feature vector is stored and also converted into a corresponding BC and stored. This results in each database having a corresponding BC copy. This will allow us to compare the underlying system accuracy to the accuracy of the system with the BC scheme by comparing feature distances and BC distances.

III. Experiment

To quantify the BC scheme's effect on the underlying system accuracy we perform a Euclidean distance test.

For each database we first run each image through all our proposed preprocessing and normalization steps. Some of the databases contained images where facial regions could not be detected and cropped from the image. We noticed that these images containing facial regions that could not be detected by the Haar Cascade Classifiers (but can be easily detected by the human eye) contain partial face images. Within these partial face images, sections of the subject's face have been covered by some object or are being cut out of the picture. We decided to remove the subjects these images belonged to from our experiment. We made this descision because our system does not allow for the registration of images in which a face region cannot be detected. We removed 11 subjects from the AT&T Face Database, 9 subjects from the Faces95 Database, and 1 subject from the BioID Face Database (see Fig. 6 for examples of images in which our proposed scheme could not detect a face).

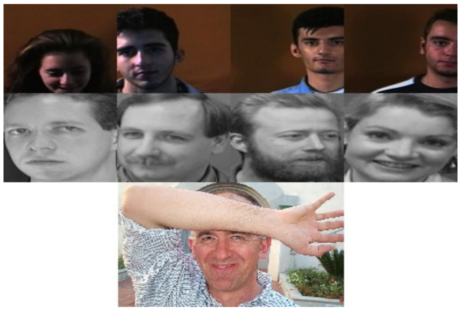

Fig. 6. Example Pictures where Face could not be Detected

After each image is preprocessed and normalized, we perform feature extraction on each image. The resulting features are stored for testing. We also convert and store the features corresponding BCs for testing.

To perform our Euclidean distance test, we define two types of data values: User values and Impostor values. A User value is the Euclidean distance between two feature vectors, or two BCs, belonging to the same subject. Correspondingly, an Imposter value is the Euclidean distance taken between two feature vectors, or two BCs, belonging to different subjects.

We perform every possible comparison between feature vectors, and BCs, of the same subject to gather all possible User values. Next, we gather an equal amount of Imposter values by comparing feature vectors, and BCs, of different subjects.

Finally we plot our results for each database (see Fig. 7). The x-axes correspond to the distance between two feature vectors. The y-axes correspond to the distance between two BCs. A single point on the graph represents the distance between two features and their corresponding BCs.

TABLE II
Database FARs and FRRs

Database	Feature FAR	Feature FRR	BC FAR	BC FRR
JAFFE	1%	26.63158%	1%	28.31579%
AT&T	1%	15.70881%	1%	18.08429%
Yale	1%	21.91571%	1%	26.59004%
Faces95	1%	6.107769%	1%	8.522974%
BioID	1%	5.137102%	1%	5.548408%
IUPUI	0%	4.07407%	0%	6.172384%

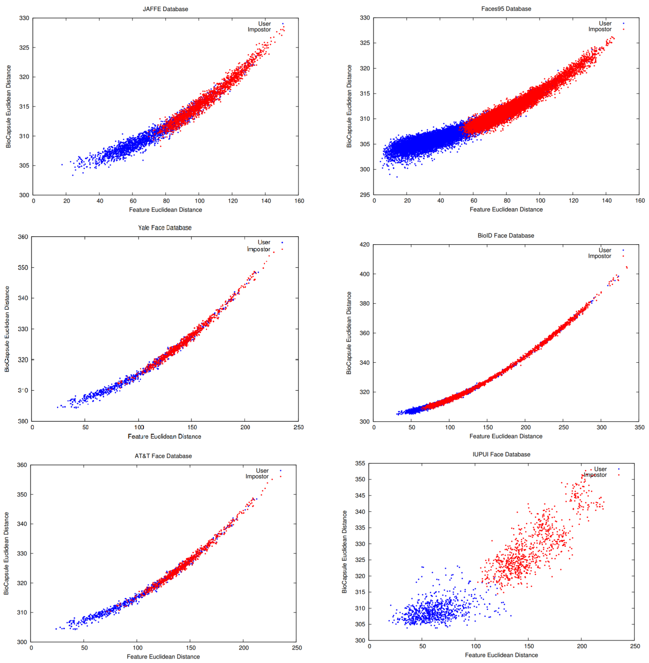

Fig. 7. Euclidean Distance Experiment Results

We present the data in this manner to demonstrate the relationship between our underlying authentication scheme and the BC embedded scheme. In order to reject impostors and accept users a threshold must be placed. For the underlying system, a vertical threshold must be placed to separate the comparisons between feature vectors that should be authenticated or rejected by the system. Likewise, a horizontal threshold must be placed to separate the comparisons between the features' corresponding BCs that should be authenticated or rejected by the BC scheme embedded system.

From our experiments we find that, by embedding the BC scheme into our system, there is not a great trade-off in system accuracy for the BC scheme's security benefits (see Table 2).

Unfortunately, for the more constrained databases, our underlying system accuracy could use improvement. We place thresholds such that the False Acceptance Rate (FAR) of the system is either 1% or 0%. The more constrained databases (JAFFE, AT&T, and Yale) suffer high a False Rejection Rate (FRR) due to the low FAR. These FRRs will hurt our system's usability in these more constrained environments.

Though the system accuracy varies between the databases, the embedded BC system accuracy always remains close to the underlying system's accuracy, and, in some cases, the BC embedded system even outperforms the underlying system. This correlation is very promising. We hope to improve our underlying system accuracy and in turn improve the BC embedded system accuracy in future research.

IV. Conclusion

Our results have shown that the BC scheme does not have a great effect on the underlying accuracy of a facial biometric authentication system. Therefore, our goal of offering cancellability and preserving privacy, while not greatly diminishing system accuracy, is achieved. Our facial authentication system performed very well in unconstrained environments. This is a good indication that an implementation of our facial authentication scheme can be extended to domains other than static authentication on a laptop.

V. Future Work

We first would like to implement a more sophisticated feature extraction scheme such as the ”Bio-inspired Features” feature extraction method demonstrated in [4] and [5]. We predict that implementing this feature extraction scheme will improve our system's accuracy when confronted with both constrained and unconstrained environments.

After improving our underlying system's feature extraction, we would also would like to implement a partial facial detection algorithm as part of our preprocessing, such as the scheme presented in [6]. Enabling our BC facial authentication scheme's implementation to detect partial faces within images will enhance its usability.

We would like to make our current facial authentication implementation able to support continuous authentication. We also would like to extend the BC scheme to other biometric modalities. Our research goal is to implement the BC scheme within an accurate continuous and multi-modal biometric scheme such as [6][7][8].

Acknowledgment

We extend our gratitude to Amy, Brandon, Carol, Claire, Darshan, Emily, Ivana, Joevita, Kaitlin, Kelby, Mark, Nash, Ravynne, Snow, Susie, Tom, and Will for allowing us to take and use their photos in our IUPUI Face Database and experiment.

References

[1] Y. Sui, X. Zou, E.Y. Du and F. Li “Design and analysis of a highly user-friendly, secure, privacy-preserving, and revocable authentication method”, IEEE Transactions on Computers, vol. 63, no.4, pp. 902-916, April 2014

[2] Xiaoyang Tan, Triggs. B, Enhanced Local Texture Feature Sets for Face Recognition Under Difficult Lighting Conditions[J], IEEE Transactions on Image Processing, 2010, 19(6): 1635-1650.

[3] V. Struc and N. Pavesic, ”Gabor-Based Kernel Partial-Least-Squares Discrimination Features for Face Recognition”, Informatica (Vilnius), vol. 20, no. 1, pp. 115–138, 2009.

[4] Guo, Guodong, et al. ”Human age estimation using bio-inspired features.” Computer Vision and Pattern Recognition, 2009. CVPR 2009.

[5] Spizhevoi, A. S., and A. V. Bovyrin. ”Estimating human age using bio-inspired features and the ranking method.” Pattern Recognition and Image Analysis 25.3 (2015): 547-552.

[6] U. Mahbub, V. M. Patel, D. Chandra, B. Barbello, and R. Chellappa. Partial face detection for continuous authentication. In 2016 IEEE International Conference on Image Processing (ICIP), pages 2991– 2995, Sept 2016

[7] Hataichanok Saevanee, Nathan Clarke, Steven Furnell, and Valerio Biscione. 2015. Continuous user authentication using multimodal biometrics. Comput. Secur. 53, C (September 2015), 234-246. DOI=http://dx.doi.org/10.1016/j.cose.2015.06.001

[8] K. Niinuma, U. Park, and A. K. Jain, ”Soft biometric traits for continuous user authentication,” Trans. Inf. Forensic Secur., vol. 5, pp. 771-780, 2010.