PRIVACY-PRESERVING FACIAL RECOGNITION USING BIOMETRIC-CAPSULES

A Thesis

Submitted to the Faculty

of

Purdue University

by

Tyler S. Phillips

In Partial Fulfillment of the

Requirements for the Degree

of

Master of Science

May 2020

Purdue University

Indianapolis, Indiana

THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF THESIS APPROVAL

Dr. Xukai Zou, Chair

Department of Computer and Information Science

Dr. Mohammad Al Hasan

Department of Computer and Information Science

Dr. Feng Li

Department of Computer and Information Technology

Approved by:

Dr. Mihran Tuceryan

Chair of Department of Computer and Information Science

To my parents, Paul and Amy.

ACKNOWLEDGMENTS

I would like to thank my advisor, Professor Xukai Zou, as well as my other committee members, Professor Mohammad Al Hasan and Professor Feng Li, for their guidance in completing this work and their mentorship throughout my studies.

TABLE OF CONTENTS

Page

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

Phillips, Tyler S. M.S., Purdue University, May 2020. Privacy-Preserving Facial Recognition Using Biometric-Capsules. Major Professor: Xukai Zou.

In recent years, developers have used the proliferation of biometric sensors in smart devices, along with recent advances in deep learning, to implement an array of biometrics-based recognition systems. Though these systems demonstrate remarkable performance and have seen wide acceptance, they present unique and pressing security and privacy concerns. One proposed method which addresses these concerns is the elegant, fusion-based Biometric-Capsule (BC) scheme. The BC scheme is provably secure, privacy-preserving, cancellable and interoperable in its secure feature fusion design.

In this work, we demonstrate that the BC scheme is uniquely fit to secure state-of-the-art facial verification, authentication and identification systems. We compare the performance of unsecured, underlying biometrics systems to the performance of the BC-embedded systems in order to directly demonstrate the minimal effects of the privacy-preserving BC scheme on underlying system performance. Notably, we demonstrate that, when seamlessly embedded into a state-of-the-art FaceNet and ArcFace verification systems which achieve accuracies of 97.18% and 99.75% on the benchmark LFW dataset, the BC-embedded systems are able to achieve accuracies of 95.13% and 99.13% respectively. Furthermore, we also demonstrate that the BC scheme outperforms or performs as well as several other proposed secure biometric methods.

1. INTRODUCTION

Through the use of biometric systems, users are able to utilize their intrinsic physiological (face, iris, fingerprint, etc.) and behavioral (speech, gait, mobile swiping patterns, etc.) biological traits in order to be recognized [1]. This grants the user the convenience of not needing to carry with them a traditional knowledge-based or physical object-based credentials (i.e. passwords or smart cards). Though biometrics-based systems offer this convenience, they also present their own set of pressing security and privacy concerns [2]. If an attacker is able to steal the biometric template of a victim, the victim's biometrics are forever lost to the attacker. The victim cannot reasonably revoke and reset their physiological or behavioral traits, as they could for a stolen password or smart card. Furthermore, through analysis of a stolen biometric template, an attacker may be able to derive private, personal information about the victim user, such as ethnicity, age, gender, health condition [3–5].

In paper [6], the authors propose the Biometric-Capsule (BC) scheme in order to address these pressing security and privacy concerns. This fusion-based cancellable biometric scheme involves the introduction of a reference subject (RS). Each user chooses (or is assigned) an RS during enrollment. Then, in order to carry out any biometric recognition task, a user's sampled biometrics are securely fused with the biometrics of their corresponding RS in order to yield a resulting BC. Through the BC scheme's secure fusion process, the contributions of the user and RS features toward the resulting BC are masked. Therefore, analysis of the resulting BC does not reveal the user or RS biometric features, even in the case most favorable to an adversary.

In this work, we embed the BC scheme into state-of-the-art facial recognition systems which leverage recently proposed deep learning-based techniques. This allows us to demonstrate several highly advantageous properties of the BC scheme and make several novel contributions:

(1) The BC scheme is interoperable in design and requires no fixed biometric sampling, detection, alignment, segmentation, feature extraction, feature representation or classification techniques in order to accommodate it. Therefore, the BC scheme can be seamlessly embedded into existing systems which use the most current and robust techniques as they are developed. This use of state-of-the-art techniques in conjunction with the BC scheme addresses several performance and flexibility issues challenging other proposed secure biometrics methods [7, 8].

(2) Through comparison of underlying systems and BC-embedded systems, we are able to directly demonstrate the performance effects of embedding the BC scheme into an underlying system. This minimal effect (and sometimes improvement) upon underlying performance provides strong motivation for the use of the BC scheme to secure state-of-the-art systems.

(3) As the BC scheme is both provably privacy-preserving and uniquely fit to secure state-of-the-art systems, it is able to effectively address emerging user concerns surrounding biometric technologies [9–12].

2. RELATED WORKS

In recent years, many methods have been proposed and investigated in hopes of robustly securing biometric templates. According to Jain et. al. [13], an ideal secure biometric system should possess many attributes including: biometric template security, cross-matching resistance, privacy-preservation and minimal negative effects on classification performance. Two broad classes of approaches for securing biometric templates have emerged: biometric cryptosystems (BCS) and cancellable biometrics (CB).

2.1 Biometric Cryptosystems

BCS approaches either bind information with biometric templates or use biometric templates directly to generate keys which are then used in place of biometric templates. Both types of approaches yield biometric-dependant public data known as helper data. This helper data is stored by the system during enrollment and, as a result, must preserve user privacy. Based on how this helper data is used within the system, BCS can be split into two sub-classes of approaches: key binding and key generation schemes. In key binding schemes, a user must provide secret information which is combined with their biometric template in order to generate helper data. Keys can then be derived from the resulting helper data. Fuzzy vault and fuzzy commitment schemes, such as [14–16], are examples of key binding schemes. In key generation schemes, helper data is derived directly from the original biometric template. As in key binding schemes, keys are derived from the resulting helper data. Fuzzy extractor and secure sketch schemes, such as [17–19], are examples of key generating schemes.

2.2 Cancellable Biometrics

CB approaches involve applying transformations directly to a biometric template such that retrieving the original biometric template is computationally costly. The altered biometric templates are then used for recognition. Then, if such a cancellable template is stolen, the attacker cannot derive the personal information of the user. In addition, the user can revoke, or cancel, the cancellable biometric template and alter their biometrics differently for future recognition tasks. CB approaches can be divided into two sub-classes: salting schemes and noninvertible transformations. In salting schemes, users provide secret information such as a password or PIN. Their biometric template is then transformed by an invertible function with respect to the provided secret information. Since these transformations are typically invertible to some extent, the secure storage of each user's corresponding secret information becomes of the utmost importance. Examples of salting schemes include [20–22]. In noninvertible transformations schemes, a biometric template is transformed using a noninvertible (or one-way) function. Unfortunately, many noninvertible transformations systems are not provably secure, and are indeed invertible under certain conditions [7]. Along with the BC scheme [6], examples of noninvertible transformations systems include [23, 24]. For both salting and noninvertible transformations schemes, the transformations applied to biometric templates must be chosen with care. On one hand, the transformations must conceal user biometrics if transformed templates are compromised. Furthermore, the transformations must preserve user privacy. On the other hand, if these transformations raise inter-class similarity or raise intra-class variability, the performance of the biometric recognition system will suffer [1, 8].

For an extensive overview of proposed BCS and CB, their respective vulnerabilities and benchmark results please see [7, 8, 13, 25].

3. BIOMETRIC-CAPSULE SCHEME

3.1 Biometric-Capsule Generation

The BC scheme is an elegant feature fusion-based CB method. Its secure fusion process involves three main steps which take place after feature extraction/representation and before classification within a biometric recognition work-flow. As shown in Fig. 3.1, the BC scheme takes two feature vectors as input, one belonging to a user and the second belonging to the user's corresponding RS. Using the two input feature vectors, three steps are carried out in order to generate a BC: (1) signature extraction, (2) key generation and (3) secure fusion. The overall BC generation workflow can be seen in Fig. 3.1. Furthermore, the Python pseudo-code for each BC generation step can be seen in Algo. 1. Here, we discuss each of these BC generation steps in detail.

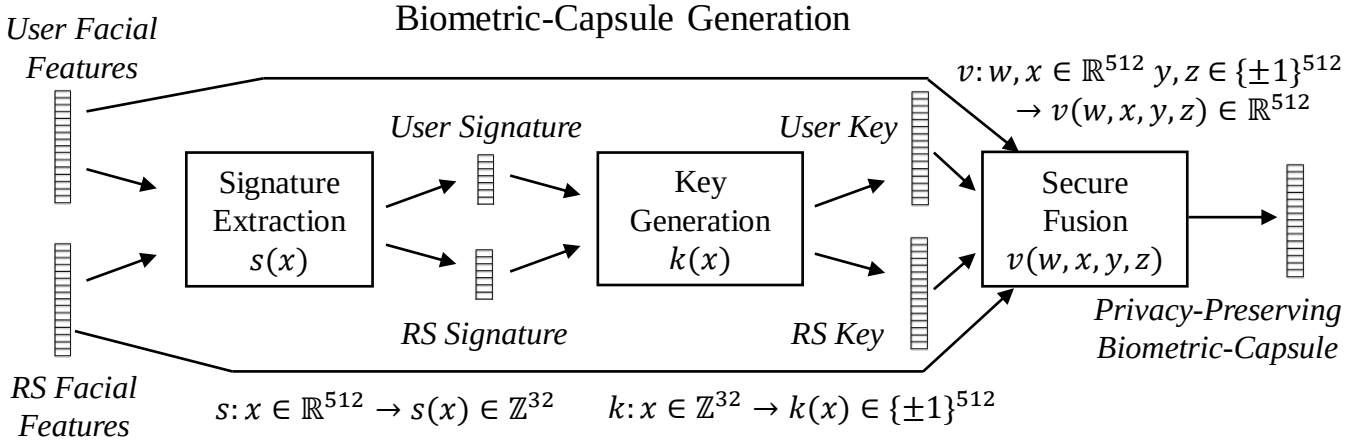

Fig. 3.1. Biometric-Capsule (BC) generation involving signature extraction, key generation and secure fusion

3.1.1 Signature Extraction

The first step in the BC generation process is signature extraction. This step involves extracting a lower-dimensional, representative signature from a facial feature vector. To perform signature extraction, we use a three-level averaging method similar to the one proposed by [26], although a different signature extraction method could be used if the system designer wishes. In our experiments, the chosen method first involves the reshaping of input feature vector from shape $\mathbb{R}^{512}$ to $\mathbb{R}^{32x16}$. Next, a $\mathbb{R}^{5x5}$ kernel is used to perform a padded averaging convolution. This means that an average value of the area covered by the kernel is found and the feature matrix is padded such that the result of the convolution is the same size as the input. Next, the difference between the original feature matrix and the resulting convolved matrix is found. Then, a row-wise average of the resulting matrix is obtained. Finally, the resulting vector values are multiplied by 10^3, rounded to integer values and mapped to values positive integer values through an absolute value operation to obtain the input feature's $\mathbb{Z}^{32}$ signature vector.

It can be seen that this signature extraction step represents a one-way function as, given an input feature vector, $x \in \mathbb{R}^{512}$, it easy to compute (as shown in Algo. 1) a signature vector, $s(x) \in \mathbb{Z}^{32}$, but, given a resulting signature vector, $s(x) \in \mathbb{Z}^{32}$, it is impossible to determine the feature vector, $x \in \mathbb{R}^{512}$, from which the signature vector was derived.

3.1.2 Key Generation

The second step of the BC generation process is key generation. The key generation process utilizes a feature's extracted signature, $x \in \mathbb{Z}^{32}$, as input. Each of a signature's 32 integer values are used as seeds in a random number generator (RNG) to generate 16 uniformly random values (for a total of 512 random values) between 0 and 1. These randomly generated values are placed into a key vector 512 (the same shape as the initial FaceNet feature embedding). Finally, all values within the key

are rounded and all resulting 0 values are changed to -1. As shown in Fig. 3.1 and Algo. 1, the key generation process will finally result in a key vector of 512 values of 1 or -1, i.e. $k(x) \in \{\pm 1\}^{512}$.

It can be seen that this key extraction process does not reveal information regarding the original input feature vector. This key extraction process represents another one-way function. Given signature value seeds, $x \in \mathbb{Z}^{32}$, and a RNG, it is easy to generate a set of randomly uniform values and then map them to a vector of 1 or -1, $k(x) \in \{\pm 1\}^{512}$, but, given a set of randomly uniform values mapped to values of 1 or -1, $k(x) \in \{\pm 1\}^{512}$, it is impossible to deterministically derive the signature value seeds, $x \in \mathbb{Z}^{32}$, used by the RNG to generate the set.

3.1.3 Secure Fusion

The final step of the BC generation process is the secure fusion step. From this secure fusion step, a resulting Biometric-Capsule is obtained. Secure fusion takes two feature vectors as input, one belonging to a user and the other belonging to the user's corresponding RS. The two keys generated using the two features are also used as inputs. The user key is used to transform the RS feature through element-wise multiplication. Likewise, the RS key is used to transform the user feature through element-wise multiplication. Through these transformations, the contribution of the features to the final resulting BC is masked. Finally, the altered biometrics are fused through an unweighted addition operation to obtain a BC (as shown in Fig. 3.1 and Algo. 1).

This Secure Fusion step can be simply represented using the following equation:

$$v(w, x, y, z) = w * z + x * y \tag{3.1}$$

where $w \in \mathbb{R}^{512}$ and $x \in \mathbb{R}^{512}$ are the user and RS features respectively, $y \in \{\pm 1\}^{512}$ and $z \in \{\pm 1\}^{512}$ are the user and RS keys respectively, $*$ is an element-wise multiplication, $+$ is a simple vector addition and $v(w, x, y, z) \in \mathbb{R}^{512}$ is the resulting BC.

A few points of this secure fusion process should be noted. First, no feature information is lost when a feature is altered through the element-wise multiplication with a key. Since the keys used for feature alteration only contain values of 1 or -1, feature values can only possibly be unaffected or negated. Second, no weight is given to the user or RS features when fusion occurs. This means the altered user and RS features contribute equally to the final, resulting BC. After a BC is generated it can be used for any biometric recognition task.

3.2 Notable Attributes of the Biometric-Capsule Scheme

In a BC-embedded biometric recognition system, the BC scheme is used to alter all biometrics sampled by the system. A BC-embedded system performs BC fusion between feature extraction/representation and classification steps (as shown in Fig. 4.1). Therefore, each time the user's biometrics are sampled by the system, the user's biometric features are fused with the biometric features of the user's corresponding RS. As a result, BCs, rather than the original user biometric features, are used for recognition tasks. Then, if an attacker infiltrates the BC-embedded system, BCs are compromised rather than unsecured and sensitive information divulging biometric templates. Furthermore, if any security concern exists, users can revoke compromised BCs and can use a different RS for BC generation in the future. In previous works, [6] indicated that the BC approach has minor negative affects on underlying iris authentication system accuracy. Furthermore, [27, 28] demonstrated that the BC scheme could be used to secure facial authentication systems and be used alongside deep learning techniques.

The elegant, simple design of the BC scheme yields highly advantageous properties. Rather than dictating which biometric sampling, segmentation, alignment, feature extraction, feature representation or classification steps occur within a biometric system in order to accommodate it, the BC scheme's flexible design allows it to instead be embedded within existing systems. This gives system designers the

Algorithm 1: Biometric-Capsule Generation Python Pseudo-Code

```
import numpy as np
from scipy.signal import convolve2d
signature_extraction (feature ∈ ℝ^512)
    lvl1 = convolve2d(
     feature.reshape(32, 16), np.ones((5, 5))/25.,
     mode =′ same′, boundary =′ wrap′)
    lvl2 = feature.reshape(32, 16) − lvl1
    lvl2 = np.average(lvl2, axis = 1) ∗ 1000.
    signature = np.around(lvl2).astype(int)
    signature = np.abs(signature)
    return signature ∈ ℤ^32
key_generation (signature ∈ ℤ^32)
    key = np.empty((0, ))
    for s in signature do
        np.random.seed(s)
        key = np.append(
         key, np.random.choice(2, 16))
    end
    key = (key ∗ 2) − 1
    return key ∈ {±1}^512
secure_fusion
  (u_feature ∈ ℝ^512, u_key ∈ {±1}^512, rs_feature ∈ ℝ^512, rs_key ∈ {±1}^512)
    bio_capsule = u_feature ∗ rs_key + rs_feature ∗ u_key
    return bio_capsule ∈ ℝ^512
```

flexibility to design an underlying biometric system how they wish, with no direct consideration for the BC scheme. After designing an underlying system, the system

designer can then embed the BC scheme within their system's pipeline between feature extraction/representation and classification steps in order to secure it and to protect the privacy of its users.

This advantageous property differs from many other secure biometric methods which make explicit or implicit constraints upon the work-flow of an underlying system in order to accommodate them. In principle, the BC scheme can be embedded into any existing biometric system and utilize the system's current biometric techniques. As shown in Fig. 3.1, only the BC scheme's secure fusion process, involving signature extraction, key generation and feature fusion steps (which themselves are flexible), must be embedded into the existing system. This allows for the BC scheme to be embedded into underlying biometric systems which use the most current and robust biometric techniques as they are developed. This is quite advantageous indeed, as many recently proposed deep learning-based biometric techniques have been shown to be extremely robust and accurate [29, 30]. Therefore, the BC scheme can leverage the highly discriminative features of underlying state-of-the-art systems, while providing robust security and privacy benefits, and only degrading underlying system performance slightly. In Fig. 3.2, we illustrate the t-SNE projection [31] of FaceNet features [32] (the most widely accepted deep learning facial feature representation method), ArcFace features [33] (the current state-of-the-art facial feature representation method) and corresponding BCs (all generated using a single RS) of 6 similar subjects with a few hundred images each. Note that, while a human may find it hard to discern inter-class differences between the classes, the FaceNet and ArcFace feature representations are able to be easily separated into distinct classes with few errors. Furthermore, for the most part, the BC versions of the underlying features preserve the clear separability of the classes while rearranging their relative positions. Due to its elegant design, the BC scheme is uniquely fit to secure biometric recognition systems which utilize these techniques.

Though the BC scheme introduces no constraints upon a system's preprocessing, feature extraction, or classification steps, the BC scheme does require the introduction

of RSs. Anytime a biometric template must be generated by the system, an RS must be retrieved by (or provided to) the system. This is because, without an RS, BC fusion is not possible.

Fortunately, how RSs are incorporated within a BC-embedded system is quite flexible. During enrollment, a new user is assigned (or chooses) a corresponding RS. The RS can be made public or kept private (with no loss in privacy benefits as shown later in this section). All users can be assigned (or choose) a unique RS, or sets of users can be assigned (or choose) the same RS. Since user and RS biometrics contribute equally in BC fusion, multiple users having the same RS introduces no security concerns as we will show later in this section. Later, when presenting their biometrics to the system in order to carry out a recognition task, the user could provide their RS to the system in a variety of different ways. A few examples are:

- In high security scenarios, an RS could be a physical object kept by the user and provided at recognition time. In this type of system, only a database of registered BCs would need to be maintained by the system. Storing RSs and information about which user(s) each RS corresponds to would not be necessary.

- A set of RSs could also be provided by the system for the user to choose from at recognition time. In this type of system, a database of registered BCs and RSs would need to be maintained by the system. Storing information about which user(s) each RS corresponds to would again not be necessary.

- The system could store and automatically use the user's corresponding RS at recognition time. This method would provide the most convenience to the user as they would not need to keep track of their RS. As a result, the BC scheme would be fully-transparent to users. Despite this transparency, users would still be protected by the BC scheme's robust security and privacy benefits. In this type of system, a database of registered BCs, RSs and information about which user(s) each RS corresponds to would need to be maintained.

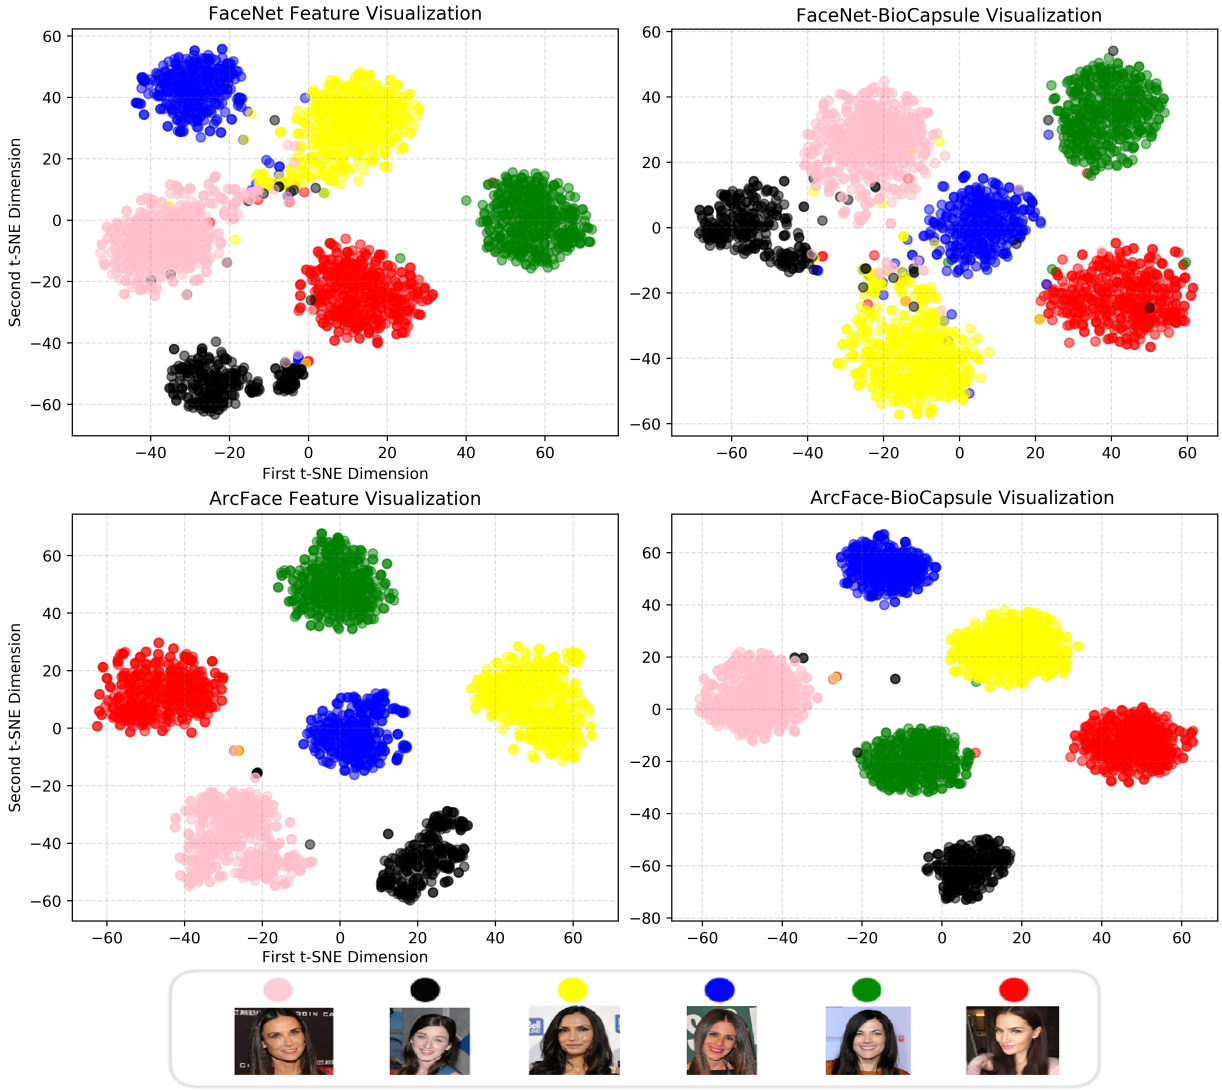

Fig. 3.2. t-SNE visualization of FaceNet, ArcFace and corresponding BC templates (using a single RS for BC generation) of 6 similar looking subjects

It should be noted that the user biometric features (which are fused with corresponding RS features to form BCs) are not stored by any of the aforementioned systems. Ultimately, how RSs are incorporated in a BC-embedded system is the system designer's choice and should reflect and enhance the use case of the underlying biometric system.

3.3 Security and Privacy of the Biometric-Capsule Scheme

In this work, we will consider a secure biometric template to be privacy-preserving if it can robustly secure and mask the user biometrics which were leveraged to generate the secure biometric template. Therefore, in this context, privacy refers to the confidentiality of unsecured biometric templates and the sensitive biometric traits that can be derived from such unsecured templates.

In addition to being flexible in design, the BC scheme also offers robust, provably secure and privacy-preserving benefits. Since the signature extraction, key generation and fusion steps of the BC scheme each have one-way properties, the resulting BC scheme can be shown to be essentially a one-way function. In paper [6], authors formally proved many security and privacy benefits of the BC scheme. These benefits include that the BC scheme is robust in defending against the following four types of attacks. (1) The first type of attack is the case in which a BC is stolen and the attacker then attempts to derive the user's biometric features, which is impossible due to that it will be equivalent to solving an underdetermined equation (as shown in Eq. 3.1). (2) The next type of attack is the case in which the attacker has stolen a user's BC and the user's corresponding RS. This will result in the attacker deriving two possibilities for each value of the user's feature vector (as they will need to guess 1 or -1 for each value within the user's key). This means that the number of possible user feature vectors will grow exponentially with respect to size of the user feature vector. In our proposed system $O(2^{512}) \approx O(10^{154})$ possible feature vectors can be derived, making obtaining the user's true feature vector computationally infeasible. (3) The third type of attack is the case in which the attacker attempts to derive the RS from multiple stolen BCs of one or multiple users, which is to solve an underdetermined system of equations and, thus, is impossible. (4) The final type of attack is the case the attacker has stolen multiple BCs (where the BCs belong to several or one user) and their corresponding RSs, which is results in many sub-cases of (2), which are computationally infeasible. The detailed, formal proofs can be found in [6].

4. DESIGN OF STATE-OF-THE-ART FACE RECOGNITION SYSTEMS

In this section we propose state-of-the-art biometric recognition systems to carry out facial verification, authentication and identification. As previously noted in Sec. 3, the BC scheme can be seamlessly embedded within these underlying systems. Since the BC scheme is flexible in design, we aimed only at using the most popular and state-of-the-art techniques (particularly involving recent deep learning techniques) while designing these underlying systems. While designing the systems, we made no direct considerations about how the chosen techniques would work in conjunction with the BC scheme.

The underlying verification, authentication and identification systems are quite similar. In fact, they perform the same steps and only differ at the classification step, as verification, authentication and identification are fundamentally different classification problems. Biometric verification is a binary classification problem in which one must determine if two biometric templates belong to the same person. Biometric authentication is a binary classification problem where one must decide if a query biometric template belongs to the enrolled subject whom it claims to be. Biometric identification is a multi-class classification problem in which one must determine the identity of a query biometric template given a group of enrolled users.

The enrollment and recognition work-flows used by the proposed BC-embedded authentication systems can be seen in Fig. 4.1. Each of the proposed systems work using three main steps: (1) preprocessing (including biometric detection, alignment and segmentation), (2) feature extraction and representation and (3) classification.

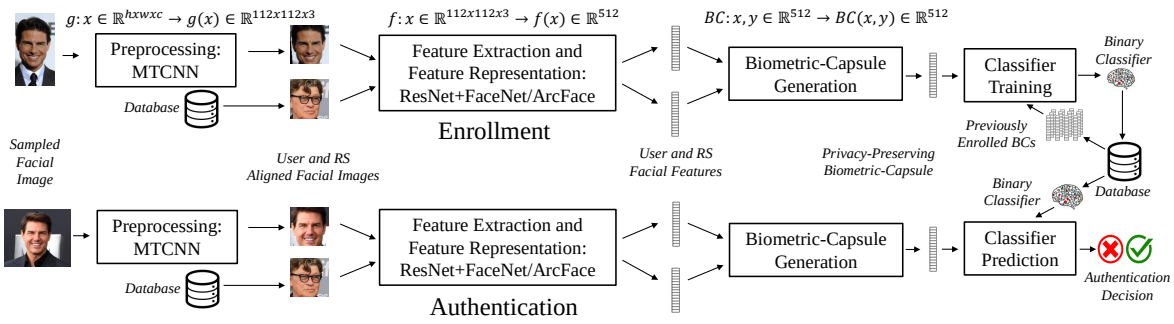

Fig. 4.1. Work-flow of the BC-embedded facial authentication system used in our experiments.

4.1 Preprocessing

The first step of each of the systems is to perform preprocessing tasks including biometric detection, alignment and segmentation. Biometric detection is the process of detecting a region of interest within a biometric signal from which features can be extracted and, in turn, can be used for recognition tasks. Alignment is then the process of normalizing the captured biometric region of interest. Finally, segmentation is the process of segmenting (in this case cropping) the relevant, aligned parts of the biometric signal for later feature extraction and recognition tasks. These tasks were accomplished through the use of detected facial bounding boxes and landmark points.

We chose to utilize the popular Multi-Task (Cascaded) Convolutional Neural Network (MTCNN) method [34] within our biometric recognition systems. This widely-accepted deep learning-based method is quite robust and out-performed other preprocessing methods which we tried. This method uses a cascade of three convolutional neural networks (CNNs) to perform detection of faces and facial landmarks. The first CNN in the cascade, the Proposal Network (P-Net), is used to generate candidate bounding boxes which potentially contain faces. The second CNN in the cascade, the Refinement Network (R-Net), takes the P-Net candidate bounding boxes as input and attempts to reject false candidates. The final CNN in the cascade, the Output Net-

work (O-Net), takes the refined R-Net candidate bounding boxes as input, attempts to reject more false candidates and finally outputs each remaining candidate detection along with five corresponding fitted facial landmark points (left/right eyes, tip of the nose, left/right sides of the mouth). We were able to find and utilize an open source implementation of MTCNN given by [35] which was trained on the WIDER FACE [36] and CelebA [37] datasets.

We leveraged information about our experimental datasets in order to decide what to do in the cases of multiple or no facial detections. In the case of multiple facial detections, we decided to select the center-most face in the image as the only face to consider in further steps within the systems' work-flows. Any other facial detections were then ignored and not considered further. In the case where no faces were detected, we decided to skip alignment and segmentation steps and directly forward the entire image to the feature extraction step. These ad-hoc heuristics worked well for our experimental setting as each image in all the experimental datasets contained a face (typically the center-most face) to be used for recognition or verification tasks. In other non-experimental settings, particularly in high security scenarios, it may be more suitable to reject images which contain no or multiple detected faces.

After we retrieved the (center-most) facial detection bounding box and five corresponding facial landmark points, we performed alignment and segmentation. To perform alignment, we performed an Affine transformation such that the two eye facial landmarks would be appear at a fixed location. We chose to include a 42-pixel margin around the bounding box in order to capture additional facial features such as chin shape, hair line, color and style, ears, etc.

4.2 Feature Extraction and Representation

The second step of the biometric recognition systems is feature extraction and representation. We decided to implement two versions of each system, each of which uses

a popular facial feature representation method, to demonstrate the interoperability of the BC scheme.

4.2.1 FaceNet Feature Representation

The first version of the systems utilizes the very popular and most widely adopted 2015 FaceNet method [32]. The system first extracts facial features from a preprocessed facial image using a deep Inception-ResNet-v1 architecture CNN [38]. Next, using the FaceNet Triplet Loss [32], extracted features are then embedded into a compact 512-dimensional space in which Euclidean distance directly corresponds to facial dissimilarty (i.e. a larger distance between feature vectors directly denotes larger facial dissimilarity).

We were able to find and utilize an open source FaceNet model given by [39] which was trained on the CASIA WebFace dataset [40].

4.2.2 ArcFace Feature Representation

The second version of the systems utilizes the current state-of-the-art feature representation method, the 2019 ArcFace method [33]. This version of the system extracts facial features from a preprocessed facial image using an extremely deep 100-layer ResNet model [41]. Next, using the ArcFace Additive Angular Margin Loss [33], features are then embedded onto a sphere where angular distance directly corresponds to facial dissimilarty.

The Additive Angular Margin Loss (AAML) gives the ArcFace method several notable advantages over the FaceNet method's Triplet Loss (TL). AAML is derived by performing slight modifications to plain Cross-Entropy and Softmax Loss, making it much more computationally efficient than TL which requires building triplets of feature embeddings before TL can be computed. Furthermore, AAML contains a margin penalty term which penalizes the correct classifications when training a model using AAML. Through the use of this margin penalty term, AAML is able to yield

discriminative features vectors in which the features of different subjects are separated by easily-interpretable linear margins. For additional details regarding the two feature representation methods please see the original papers, [32,33], or deep face recognition survey [30].

We were able to find and utilize an open source ArcFace model given by the ArcFace authors [35] which was trained on the MS-Celeb-1M dataset [42].

It should be noted that both the FaceNet or ArcFace feature extraction and representation models yield 512-dimensional feature vectors. These feature vectors can be used in BC generation as shown in Sec. 3 in a BC-embedded facial recognition system or directly for recognition tasks in an unsecured, underlying system.

It should also be noted that the inversion of facial feature vectors, such as the feature vectors produced by the FaceNet and ArcFace method, to their corresponding facial images is an active area of research. In recent works [43, 44], researchers have proposed effective methods which transform facial feature embeddings to facial images that visually reveal the private personal information (ethnicity, gender, age, etc.) of users. Such inversions would not be applicable to fused BC templates. Therefore, BC fusion is able to preserve user privacy while, at the same time, make use of effective deep learning feature embedding techniques, such as the FaceNet and ArcFace methods.

4.3 Classification

After extracting features from a preprocessed image, the BC-embedded system will then carry out BC generation using the steps outlined in Sec. 3. After BC generation, the BC-embedded system is ready to perform classification and carry out recognition tasks. An unsecured underlying system, on the other hand, will be ready to perform classification directly after performing feature extraction and representation. In either case, classification will work exactly the same as the unsecured feature templates and secured BC templates are 512-dimensional vectors.

To perform verification, we simply compare the extracted features/BCs of two facial images. We obtain the Euclidean distance between the features/BCs. If this distance is greater than a predefined threshold (obtained through analysis of training comparisons), the images are predicted as belonging to different subjects. Likewise, if the distance between the features/BCs is equal to or below the predefined threshold, the images are predicted as belonging to the same subject.

To perform authentication, a binary Logistic Regression (LR) classifier is trained for each subject during enrollment. Each subject's LR is trained with using all enrolled features/BCs. The features/BCs of the LR's corresponding subject are used as positive samples. Every other subject's enrolled features/BCs are then used as negative samples. Given a query feature/BC and a subject the query feature/BC claims to be, the authentication system classifies the query feature/BC using the claimed subject's binary LR. This results in a binary classification decision indicating whether the query feature/BC is predicted to be the subject whom they claim to be. If the classifier indicates the test feature/BC is the subject, the feature/BC is authenticated by the system (or rejected otherwise).

To perform identification, a single multiclass LR model is used. All registered features/BCs in the biometric identification system are given to the multiclass LR for training. Query features/BCs are classified by the multiclass LR as the subject whose registered features/BCs most closely match the query feature/BC.

5. EXPERIMENT

We begin our experiment by comparing the performance of the proposed underlying authentication, identification and verification systems with the performance of the BC-embedded systems. These comparisons directly reveal how embedding the BC scheme into an existing biometric system will affect the underlying system's performance. Then, we also compare the BC scheme to many popular cancellable biometrics (CB) and biometric cryptosystem (BCS) approaches. These comparisons further demonstrate the novel, advantageous attributes of the BC scheme. Overall statistics regarding each of the experimental datasets can be seen in Table 5.1.

For the BC-embedded version of each system, we generate BC templates by fusing any extracted user feature with a single, shared RS. Use of a single, shared RS allows the system to automatically generate BCs without any additional input from a users (i.e. only their sampled biometrics are still needed for recognition tasks). Therefore, the BC scheme is highly usable and fully transparent to users. As the usability of the single RS BC-embedded systems is equivalent to the underlying systems, comparisons of system performances only gauge the relative effectiveness of using features verses BCs in recognition tasks. As previously stated in Sec. 3, use of a single, shared RS does not compromise the privacy of sensitive user information. It should be noted though, if a unique RS is assigned to each user, kept secret by the user and presented to the system at recognition time, the BC scheme performance is likely to outperform the underlying system at the cost of reduced usability [28].

Before analyzing the relative performance of the BC-embedded system, the overhead of the BC scheme should be noted. Using a Dell laptop's 2-core Intel Core i7-6500u CPU, generating each BC takes ˜0.012 seconds. It should be noted that the BC generation time is substantially faster than preprocessing and feature extraction/representation steps which take ˜0.2 and ˜0.185 seconds respectively using the

same Intel Core i7-6500u CPU. Therefore, the inclusion of the BC scheme does not greatly affect the scalability of an underlying biometric system in terms of efficiency. This is particularly true if the BC scheme is embedded into systems in which RS features and keys are pre-computed. Furthermore, in terms of storage, a BC is equivalent to an underlying feature embedding.

For each test, we report several metrics commonly used to evaluate biometric recognition systems [1], such as: total false positive classifications (Total FP), total fall negative classifications (Total FN), total misclassifications (ERR), accuracy (ACC), precision (PRE), recall (REC), F1-score (F1), false acceptance rate (FAR), false rejection rate (FRR), equal error rate (EER) and area under the receiver operating characteristic curve (AUC).

Table 5.1.
Experimental Datasets Overall Statistics

Dataset	Number of Subjects	Number of Images
ORL [45]	10	400
Yale Faces [46]	15	165
Yale Faces B [47]	28	16,128
IMM [48]	40	240
Caltech Faces [49]	28	445
GTDB [50]	50	750
FEI [51]	200	2,800
FERET Color [52]	994	11,338
CMU (Pose) [53]	68	12,240
CMU (Illumination) [53]	68	21,216
CMU (Expression) [53]	68	3,016
LFW [54]	13,233	5,749
LFW (Subset) [54]	423	5,985

Table 5.2.
Verification Experiment Results

Dataset	Method	Total FP	Total FN	ACC (%)	PRE (%)	REC (%)	F1 (%)
LFW [54]	FaceNet	38	131	97.1833	98.7071	95.6333	97.1344
	FaceNet+BC	118	174	95.1333	96.0076	94.2000	95.0837
	ArcFace	0	15	99.7500	100	99.5000	99.7484
	ArcFace+BC	12	40	99.1333	99.5988	98.6667	99.1256

5.1 Facial Verification

For our verification experiment, we utilize the highly unconstrained, benchmark Labeled Faces in the Wild (LFW) dataset [54]. This dataset contains 13,233 images of 5,479 subjects. We compare the performance of the underlying system and the BC-embedded system under the View 2 testing method defined for the LFW dataset [55]. The View 2 testing method provides a predefined 10-fold cross validation split of the dataset. Each testing fold contains 300 matching and 300 mismatching biometric template comparisons. Each training fold contains approximately 575 subjects and their corresponding images. These training images can be used to generate biometric templates which can then be used to generate pair-wise training comparisons. Based on the training comparisons, later test comparisons can be predicted as matching/mismatching subject template comparisons. Therefore, this verification experiment directly evaluates the discriminative power of the FaceNet, ArcFace and corresponding BC templates. For reference, DeepFace authors [56], have reported the human accuracy in LFW verification as 97.53%. The macro-average results of the 10-fold cross validation experiment can be seen in Table 5.2.

As seen in Table 5.2, the underlying FaceNet system achieves an accuracy of 97.1833%. It should be noted that this result is significantly lesser than the 99.63% accuracy reported by the FaceNet authors as they trained their model with a private dataset of over 200 million images [32], whereas our FaceNet model was trained using

a dataset of 450 thousand images [40]. The underlying ArcFace system, on the other hand, achieves an accuracy of 99.13%, surpassing human accuracy by a significant margin [56]. This accuracy is much more similar to the current state-of-the-art result, 99.82%, reported by the ArcFace authors [33].

The effect of the BC scheme on the underlying systems can also been seen in Table 5.2. As one would suspect, the BC-embedded system has lesser performance than their underlying system counterparts. One interesting observation (which the reader may notice in the following experiments as well) is that the BC-embedded ArcFace system is able to outperform the BC-embedded FaceNet system. Furthermore, the BC scheme decrements the accuracy of the underlying FaceNet system more than the underlying ArcFace system. This is likely due to the nature of the BC generation process. As shown in Algo. 1, the BC generation process derives representative keys from input user and RS feature vectors. As we use a fixed RS in all experiments, the RS feature and key used in BC generation will always be constant. This will likely lower inter-class variability and partially account for for the lesser performance of BC-embedded systems. The use of user features and keys in BC generation must also be considered. If the underlying user feature representation is able to produce better intra-class compactness and inter-class variability, more discriminative user keys will be generated during the BC generation process. Likewise, the if the underlying user feature representation produces lesser intra-class compactness and inter-class variability, less discriminative user keys will be generated during the BC generation process and, as a result, further degrade the performance of the BC-embedded system. Following this logic, it reasonable to assume the BC scheme will have a lesser negative impact on underlying systems if they use superior state-of-the-art biometrics techniques. If lesser techniques are used, the BC scheme can be assumed to have a more adverse effect on underlying performance.

Table 5.3.
Authentication Experiment Results

Dataset	Method	Total FP	Total FN	ACC (%)	FAR (%)	FRR (%)	EER (%)	AUC
ORL [45]	FaceNet	1	0	99.9938	0.0064	0	0	1
	FaceNet+BC	1	1	99.9875	0.0064	0.2500	0	1
	ArcFace	0	0	100	0	0	0	1
	ArcFace+BC	0	0	100	0	0	0	1
Yale Faces [46]	FaceNet	1	0	99.9596	0.0433	0	0	1
	FaceNet+BC	0	1	99.9596	0	0.6061	0.1299	0.9999
	ArcFace	0	1	99.9596	0	0.6061	0	1
	ArcFace+BC	0	1	99.9596	0	0.6061	0.3030	0.9999
IMM [48]	FaceNet	11	0	99.8854	0.1175	0.0000	0.0107	0.9999
	FaceNet+BC	2	11	99.8646	0.0214	4.5833	0.2350	0.9999
	ArcFace	0	0	100	0	0	0	1
	ArcFace+BC	0	0	100	0	0	0	1
Caltech Faces [49]	FaceNet	8	1	99.9222	0.0719	0.2247	0.2337	0.9996
	FaceNet+BC	4	3	99.9395	0.0360	0.6742	0.2337	0.9999
	ArcFace	1	1	99.9827	0.0090	0.2247	0.2247	0.9998
	ArcFace+BC	0	2	99.9827	0	0.4494	0.2247	0.9999
GTDB [50]	FaceNet	33	0	99.9120	0.0898	0	0.0027	1
	FaceNet+BC	10	5	99.9600	0.0272	0.6667	0.0599	0.9999
	ArcFace	0	0	100	0	0	0	1
	ArcFace+BC	0	0	100	0	0	0	1
FEI [51]	FaceNet	528	35	99.8995	0.0948	1.2500	0.8710	0.9965
	FaceNet+BC	160	88	99.9557	0.0287	3.1429	1.0606	0.9965
	ArcFace	223	26	99.9555	0.0400	0.9285	0.9549	0.9962
	ArcFace+BC	117	37	99.9725	0.0210	1.3214	0.9508	0.9954
FERET Color [52]	FaceNet	9220	501	99.9137	0.0819	4.4188	1.6111	0.9964
	FaceNet+BC	2036	1592	99.9678	0.0181	14.0414	2.0082	0.9961
	ArcFace	781	491	99.9887	0.0069	4.3306	1.6671	0.9950
	ArcFace+BC	74	1152	99.9891	0.0007	10.1607	2.1131	0.9939
LFW (Subset) [54]	FaceNet	2779	76	99.8872	0.1100	1.2698	0.4486	0.9997
	FaceNet+BC	857	607	99.9422	0.0339	10.1420	0.9943	0.9995
	ArcFace	18	8	99.9990	0.0007	0.1337	0.0671	0.9999
	ArcFace+BC	2	78	99.9968	0.0001	1.3033	0.0854	0.9998

5.2 Facial Authentication

For our authentication experiment, we performed a 5-fold cross validation experiment for several datasets. As authentication is a binary classification task performed with respect to a single enrolled subject, each query template was classified by each subject's binary classifier (trained using all training templates with respect to the given subject). Therefore, the number of total classifications for each dataset was equal to the total number of subjects times the total number of images. Furthermore, the number of genuine authentication attempts for a dataset is then equal to the total number of images contained in that dataset, while the number of false authentication attempts is equal to the total number of subjects minus one times the total number of images. The results of the authentication experiment are shown in Table 5.3. The results shown were acquired by taking the micro-average of each subject's classification results for a given fold and finally macro-averaging the results of the five folds.

As shown in the results, the observations made in the previous verification experiment hold true. Across all experiments, the underlying ArcFace system out performs the underlying FaceNet system in terms of equal error rate. The BC-embedded systems then raise the equal error rate of the underlying version of the systems. In general, the FaceNet system accuracy is more adversely affected by the inclusion of the BC scheme than the ArcFace system.

A few interesting observations should be noted. In all cases in which the ArcFace system achieves perfect performance, the BC-embedded ArcFace system is also able to achieve perfect accuracy. Furthermore, in many cases, the BC-embedded systems make less total false positive classifications than their underlying system counterparts, despite the BC-embedded systems having a greater equal error rate. This is quite noteworthy as in most real-world scenarios false positive classifications are much more concerning from a security and privacy standpoint than false rejections. Also, a higher false rejection rate is a problem which can be easily remedied in a facial authentication system as fast re-authentication attempts can automatically be made.

Table 5.4.
Identification Experiment Results

Dataset	Method	Total ERR	ACC (%)	PRE (%)	REC (%)	F1 (%)
Yale Faces B [47]	FaceNet	2908	81.9693	82.6303	81.9704	82.0075
	FaceNet+BC	3364	79.1419	79.5990	79.1424	79.1773
	ArcFace	1519	90.5816	90.9412	90.5812	90.5755
	ArcFace+BC	2069	87.1714	87.4307	87.1719	87.1496
FEI [51]	FaceNet	33	98.8214	99.1879	98.8333	98.8190
	FaceNet+BC	47	98.3214	98.7429	98.3667	98.3005
	ArcFace	26	99.0714	99.4690	99.0833	99.0998
	ArcFace+BC	29	98.9643	99.2700	98.9667	98.9462
FERET Color [52]	FaceNet	997	91.2065	83.7165	86.2266	83.8239
	FaceNet+BC	1461	87.1141	75.7137	79.4425	76.1193
	ArcFace	423	96.2691	93.7746	94.5567	93.5721
	ArcFace+BC	603	94.6816	90.1699	91.7753	90.1506
CMU (Pose) [53]	FaceNet	8	99.9346	99.9366	99.9346	99.9346
	FaceNet+BC	11	99.9101	99.9146	99.9101	99.9102
	ArcFace	4	99.9673	99.9682	99.9673	99.9672
	ArcFace+BC	5	99.9592	99.9603	99.9591	99.9590
CMU (Illumination) [53]	FaceNet	7422	65.0170	71.7019	65.0158	66.6511
	FaceNet+BC	7651	63.9376	66.8468	63.9350	64.6440
	ArcFace	6323	70.1971	82.1605	70.1973	73.8058
	ArcFace+BC	6816	67.8733	71.9017	67.8732	68.9093
CMU (Expression) [53]	FaceNet	21	99.3037	99.3956	99.3317	99.3195
	FaceNet+BC	33	98.9057	99.0388	98.9300	98.9143
	ArcFace	7	99.7679	99.8436	99.8048	99.8115
	ArcFace+BC	6	99.8011	99.8419	99.8316	99.8277
LFW (Subset) [54]	FaceNet	151	97.4770	95.0998	96.3447	95.1583
	FaceNet+BC	316	94.7201	88.2921	90.2411	88.4626
	ArcFace	4	99.9332	99.9492	99.9632	99.9482
	ArcFace+BC	6	99.8997	99.8872	99.8923	99.8784

5.3 Facial Identification

For our identification experiment, we also performed a 5-fold cross validation experiment. As facial identification involves a single multi-class classifier used for all classification tasks, the total number of classifications performed for each dataset is simply equal to the number of total images in the dataset. To report the experimental results in Table 5.4, we micro-average the classification results with respect to each subject using a one-vs-all approach and finally macro-average results of each of the 5 cross-validation folds.

The identification results follow the same overall trends as verification and authentication experiments. The underlying ArcFace system is always able to outperform the underlying FaceNet system. In general, the BC-embedded systems diminish underlying system performance in such a way that is proportional to underlying system performance. In the case CMU (Expression) dataset [53], the BC-embedded ArcFace system actually outperforms the underlying ArcFace system, albeit by only a single less misclassification.

5.4 Comparison with Existing Methods

We also compared the BC scheme with many popular biometric cryptosystem (BCS) and cancellable biometrics (CB) methods. We test the proposed ArcFace BC-embedded system using the same dataset and testing method of several popular secure biometric methods. This allows us to compare the BC scheme's performance against other proposed secure biometric schemes. The results of each of these comparisons can be seen in Table 5.5. We report our results in terms of the metric(s) used in the original paper of the technique which we compare the BC scheme with.

The first method CB method we compared the BC scheme to was the minimum average correlation energy (MACE) cancellable filtering based method [20]. This method requires users to provides both their facial biometrics and a secret PIN during enrollment. The provided PIN is used as a seed in order to generate a random

filter. The random filter is then used to filter and encrypt the user's sampled facial images. Finally, the resulting encrypted facial images are is transformed into a corresponding MACE filter and stored by the system. At authentication time, the user again provides their facial biometrics and secret PIN. The PIN is again used to generate a random filter which is applied to the user's query facial image. Finally, the user's stored MACE filter is applied to the user's filtered facial image. Following the application of the MACE filter, the authors examine the resulting peak-to-sidelobe ratio (PSR) in order to make an authentication decision. In their paper [20], the authors perform a verification experiment using the CMU (Illumination) dataset [53]. They report an EER of 0% which we were also able to achieve using both the BC-embedded ArcFace system.

The second CB method we compared the BC scheme with was the Cancellable 2DPCA method proposed by [24]. The authors use of polynomial functions and co-occurrence matrices in order to modify facial images. They then use principal component analysis (PCA) for feature extraction. The authors use the Olivetti Research Laboratory (ORL) dataset [45] (also known as the AT&T dataset) for their experiment, and report an authentication accuracy of 96%. The BC-embedded ArcFace system was able to achieve an accuracy of 100%.

The next method we compared the BC scheme to was the BCS key binding Fuzzy Vault based method for faces [16]. This method fuses the biometric template of a user with a key the user must also provide. The authors perform an authentication experiment and report a best FAR of 5.26% and a best FRR of 26%. The BC-embedded ArcFace system was able to achieve a FAR of 0% with a FRR of 0%.

Next, we compared the BC scheme to the CB Mixing Biometrics method [23]. In many respects, this method is more similar to the BC scheme than any other method which we compared the BC method with. The Mixing Biometrics method uses the facial landmarks of a user facial image and the facial landmarks of a RS-like image in order to fuse the two faces. Classification is then performed using the fused face. Though this method is similar to the BC method in some respects, the

BC method has clear advantages. The Mixing Biometrics method requires certain predefined alignment steps to take place for later facial fusing. The BC requires no fixed alignment steps. The BC approach also preserves user privacy. From a stolen BC, an attacker cannot derive personal information (such as gender, ethnicity, age, etc.) about the victim, even when the RS image is also stolen. Unfortunately, with Mixing Biometrics fused faces, it would not be difficult for attackers to derive personal information of the user. The personal information (such as gender, ethnicity, age, etc.) of the user is clearly visible in the Mixing Biometrics fused face template. Furthermore, if the attacker obtained the fused face and the RS-like image used for facial fusion in the Mixing Biometrics method, the attacker would certainly be able to derive the personal information of the victim user by reversing the facial fusion process. The Mixing Biometrics authors use the IMM face dataset [48] for an identification experiment, as the IMM face dataset has pre-annotated facial landmarks. The authors report an EER of 6%. The BC-embedded ArcFace system was able to achieve an EER of 0%.

The final method we compared the BC scheme to was the CB Secure Computation of Face Identification method (SCiFI) [57]. This method uses a secure multi-party computation of Eigenfaces [58] in order to identify faces. Authors perform an identification on the CMU (Pose) dataset [53]. The authors report a (rank one) true positive rate of 80%. The BC-embedded ArcFace system was able to achieve (rank one) true positive rates of 100%.

Each of these comparisons demonstrates that the BC is able to perform as well or outperform other proposed secure biometric methods. It should be noted that many of the compared schemes use very constrained datasets for their experiments. Had these methods reported results using an unconstrained dataset, like the LFW dataset [54], these methods' performances would likely be very poor in comparison with the BC scheme's performance. In addition to the BC's superior performance, the BC provides many additional advantages over the compared methods. For instance, the BC scheme is flexible in design unlike most of the compared methods which assume fixed

preprocessing, feature extraction and/or classification techniques. This prevents the compared methods from using state-of-the-art deep learning-based methods. The BC is also provably secure and privacy preserving unlike some of the compared methods.

Table 5.5.
Comparison of the Biometric-Capsule Scheme with Popular Methods

Method	Dataset	Domain	Metric	Result
MACE [20]	CMU PIE (Subset) [53]	Verification	EER	0%
ArcFace+BC				0%
Cancellable 2DPCA [24]	ORL [45]	Authentication	ACC	96%
ArcFace+BC				100%
Fuzzy Vault [16]	ORL [45]	Authentication	FAR, FRR	5.26%, 23%
ArcFace+BC				0%, 0%
Mixing Biometrics [23]	IMM [48]	Identification	EER	6%
ArcFace+BC				0%
SCiFI [57]	CMU PIE (Subset) [53]	Identification	TPR	80%
ArcFace+BC				100%

6. FUTURE WORK

One future work direction is investigating the intraclass and interclass similarity and variation of BCs formed by a composition of a user's biometrics fused with the biometrics of multiple RSs. In the proposed BC scheme, if the BC database is compromised, users are able to securely revoke their compromised BCs and register new ones. If users could instead fuse their compromised BCs with an additional, secondary RSs, the new BCs could then be used for future recognition tasks. In the future, users would simply need to form BCs using their biometrics and their first RS and then fuse the resulting BC with the new, secondary RS. We need to investigate the performances of systems which use such multi-RS BC compositions in order to determine their usability.

One other future work direction is to investigate the use of dynamic RSs. In the currently proposed scheme, an RS is a fixed facial image. It may, in fact, be possible to use a subject RS rather than an image as a RS. For example, rather than using a fixed image of George Washington as an RS, George Washington the subject could be used as an RS and any image of George Washington could be used for BC fusion. The usability of dynamic RSs would most likely depend on the intra-class compactness of the feature vectors extracted from the RS subject's images.

7. SUMMARY

We have shown that the BC method can be used to effectively secure biometrics systems used for for facial verification, authentication and identification. In each of these domains, the BC scheme can be embedded into existing biometric recognition systems with virtually no constraint on how the underlying system operates. This flexible design of the BC scheme allowed us to embed the BC scheme in recognition schemes which used state-of-the-art deep learning techniques. The BC scheme offers the underlying system robust security and privacy benefits while, at the same time, affecting the underlying system's performance in a predictable manner. Furthermore, we have shown that the BC system performs as well as or outperforms many popular secure biometric techniques.

8. PUBLICATIONS

In this section, a complete list of my publications at the time of writing is given (in chronological order):

1. T. Phillips, X. Zou, and F. Li, "A Cancellable and Privacy-Preserving Facial Biometric Authentication Scheme", IEEE 14th International Conference on Mobile Ad Hoc and Sensor Systems (MASS), Orlando, FL, 2017, pp. 545-549.

2. T. Phillips, K. Byrd, and X. Zou, "A New Look at Old Abe's Color Guard: Researchers Combine Classic and Cutting-Edge Techniques to Reexamine the Identities of Soldiers in an Iconic Image", Military Images Magazine, Spring Issue 2019, pp. 60-64

3. T. Phillips, X. Zou, F. Li and N. Li "Enhancing Biometric-Capsule-based Authentication and Facial Recognition via Deep Learning", In Proceedings of the 24th ACM Symposium on Access Control Models and Technologies (SACMAT '19). ACM, New York, NY, USA, 141-146

4. T. Phillips, X. Yu, B. Haakenson and X. Zou, "Design and Implementation of Privacy-Preserving, Flexible and Scalable Role-Based Hierarchical Access Control," 2019 First IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA), Los Angeles, CA, USA, 2019, pp. 46-55.

REFERENCES

REFERENCES

[1] A. K. Jain, A. Ross, and S. Prabhakar, "An introduction to biometric recognition," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 14, no. 1, pp. 4–20, Jan 2004.

[2] S. Prabhakar, S. Pankanti, and A. K. Jain, "Biometric recognition: security and privacy concerns," *IEEE Security Privacy*, vol. 99, no. 2, pp. 33–42, March 2003.

[3] A. Dantcheva, P. Elia, and A. Ross, "What else does your biometric data reveal? a survey on soft biometrics," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 3, pp. 441–467, March 2016.

[4] E. Eidinger, R. Enbar, and T. Hassner, "Age and gender estimation of unfiltered faces," *IEEE Transactions on Information Forensics and Security*, vol. 9, no. 12, pp. 2170–2179, Dec 2014.

[5] X. Geng, Z. Zhou, and K. Smith-Miles, "Automatic age estimation based on facial aging patterns," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 12, pp. 2234–2240, Dec 2007.

[6] 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.

[7] C. Rathgeb and A. Uhl, "A survey on biometric cryptosystems and cancelable biometrics." *EURASIP J. Information Security*, vol. 2011, p. 3, 2011.

[8] U. Uludag, S. Pankanti, S. Prabhakar, and A. K. Jain, "Biometric cryptosystems: issues and challenges," *Proceedings of the IEEE*, vol. 92, no. 6, pp. 948–960, June 2004.

[9] K. Conger, R. Fausset, and S. F. Kovaleski, "San francisco bans facial recognition technology," *https://www.nytimes.com/2019/05/14/us/facial-recognition-ban-san-francisco.html*, May 14, 2019. Accessed in May 2019.

[10] S. Musil, "Microsoft calls for regulation of facial-recognition technology," *https://www.cnet.com/news/microsoft-calls-for-regulation-of-facial-recognition-technology/*, December 6, 2018. Accessed in May 2019.

[11] P. Mozur, "Inside china's dystopian dreams: A.i., shame and lots of cameras," *https://www.nytimes.com/2018/07/08/business/china-surveillance-technology.html*, July 8, 2018. Accessed in May 2019.

[12] J. Kahn, "Deep learning 'godfather' bengio worries about china's use of ai," *https://www.bloomberg.com/news/articles/2019-02-02/deep-learning-godfather-bengio-worries-about-china-s-use-of-ai*, February 2, 2019. Accessed in May 2019.

[13] A. K. Jain, K. Nandakumar, and A. Nagar, "Biometric template security," *EURASIP J. Adv. Signal Process*, vol. 2008, pp. 113:1–113:17, Jan. 2008.

[14] M. Ao and S. Z. Li, "Near infrared face based biometric key binding," in *Proceedings of the Third International Conference on Advances in Biometrics*, ser. ICB '09. Berlin, Heidelberg: Springer-Verlag, 2009, pp. 376–385.

[15] H. Lu, K. Martin, F. Bui, K. N. Plataniotis, and D. Hatzinakos, "Face recognition with biometric encryption for privacy-enhancing self-exclusion," in *2009 16th International Conference on Digital Signal Processing*, July 2009, pp. 1–8.

[16] L. Wu and S. Yuan, "A face based fuzzy vault scheme for secure online authentication," in *2010 Second International Symposium on Data, Privacy, and E-Commerce*, Sept 2010, pp. 45–49.

[17] X. Boyen, "Reusable cryptographic fuzzy extractors," in *Proceedings of the 11th ACM Conference on Computer and Communications Security*, ser. CCS '04. New York, NY, USA: ACM, 2004, pp. 82–91.

[18] Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith, "Fuzzy extractors: How to generate strong keys from biometrics and other noisy data," *SIAM J. Comput.*, vol. 38, no. 1, pp. 97–139, Mar. 2008.

[19] Q. Li, Y. Sutcu, and N. Memon, "Secure sketch for biometric templates," in *Proceedings of the 12th International Conference on Theory and Application of Cryptology and Information Security*, ser. ASIACRYPT'06. Berlin, Heidelberg: Springer-Verlag, 2006, pp. 99–113.

[20] M. Savvides, B. V. K. V. Kumar, and P. K. Khosla, "Cancelable biometric filters for face recognition," in *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.*, vol. 3, Aug 2004, pp. 922–925 Vol.3.

[21] Y. Wang and K. N. Plataniotis, "Face based biometric authentication with changeable and privacy preservable templates," in *2007 Biometrics Symposium*, Sep. 2007, pp. 1–6.

[22] 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, Dec 2006.

[23] A. Ross and A. Othman, "Visual cryptography for biometric privacy," *IEEE Transactions on Information Forensics and Security*, vol. 6, no. 1, pp. 70–81, March 2011.

[24] M. A. Dabbah, W. L. Woo, and S. S. Dlay, "Secure authentication for face recognition," in *2007 IEEE Symposium on Computational Intelligence in Image and Signal Processing*, April 2007, pp. 121–126.

[25] V. M. Patel, N. K. Ratha, and R. Chellappa, "Cancelable biometrics: A review," *IEEE Signal Processing Magazine*, vol. 32, no. 5, pp. 54–65, Sept 2015.

[26] Y. Du, R. Ives, D. M. Etter, and T. Welch, "Use of one-dimensional iris signatures to rank iris pattern similarities," *Optical Engineering*, vol. 45, no. 3, Mar. 2006.

[27] T. Phillips, X. Zou, and F. Li, "A cancellable and privacy-preserving facial biometric authentication scheme," in *2017 IEEE 14th International Conference on Mobile Ad Hoc and Sensor Systems (MASS)*, Oct 2017, pp. 545–549.

[28] T. Phillips, X. Zou, F. Li, and N. Li, "Enhancing biometric-capsule-based authentication and facial recognition via deep learning," *The ACM Symposium on Access Control Models and Technologies (SACMAT'19), (accepted)*, June 2019.

[29] K. Sundararajan and D. L. Woodard, "Deep learning for biometrics: A survey," *ACM Comput. Surv.*, vol. 51, no. 3, pp. 65:1–65:34, May 2018.

[30] M. Wang and W. Deng, "Deep face recognition: A survey," 2018.

[31] L. van der Maaten and G. Hinton, "Visualizing data using t-SNE," *Journal of Machine Learning Research*, vol. 9, pp. 2579–2605, 2008. [Online]. Available: http://www.jmlr.org/papers/v9/vandermaaten08a.html

[32] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015, pp. 815–823.

[33] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[34] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE Signal Processing Letters*, vol. 23, no. 10, pp. 1499–1503, Oct 2016.

[35] deepinsight. (2017) Arcface and mtcnn github repository. Available at: https://github.com/deepinsight/insightface.

[36] S. Yang, P. Luo, C. C. Loy, and X. Tang, "Wider face: A face detection benchmark," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[37] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *Proceedings of International Conference on Computer Vision (ICCV)*, 2015.

[38] C. Szegedy, S. Ioffe, and V. Vanhoucke, "Inception-v4, inception-resnet and the impact of residual connections on learning," in *AAAI*, 2017.

[39] D. Sandberg. (2015) Facenet and mtcnn github repository. Available at: https://github.com/davidsandberg/facenet.

[40] D. Yi, Z. Lei, S. Liao, and S. Z. Li, "Learning face representation from scratch," *CoRR*, vol. abs/1411.7923, 2014.

[41] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 770–778.

[42] Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao, "Ms-celeb-1m: A dataset and benchmark for large-scale face recognition," in *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part III*, ser. Lecture Notes in Computer Science, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds., vol. 9907. Springer, 2016, pp. 87–102. [Online]. Available: https://doi.org/10.1007/978-3-319-46487-9_6

[43] A. Zhmoginov and M. B. Sandler, "Inverting face embeddings with convolutional neural networks," *CoRR*, vol. abs/1606.04189, 2016.

[44] F. Cole, D. Belanger, D. Krishnan, A. Sarna, I. Mosseri, and W. T. Freeman, "Synthesizing normalized faces from facial identity features," *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3386–3395, 2017.

[45] F. S. Samaria and A. C. Harter, "Parameterisation of a stochastic model for human face identification," in *Proceedings of 1994 IEEE Workshop on Applications of Computer Vision*, Dec 1994, pp. 138–142.

[46] A. Georghiades, P. Belhumeur, and D. Kriegman, "From few to many: Illumination cone models for face recognition under variable lighting and pose," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, pp. 643–660, 06 2001.

[47] K. chih Lee, J. Ho, and D. J. Kriegman, "Acquiring linear subspaces for face recognition under variable lighting," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, pp. 684–698, 2005.

[48] M. M. Nordstrøm, M. Larsen, J. Sierakowski, and M. B. Stegmann, "The IMM face database - an annotated dataset of 240 face images," Informatics and Mathematical Modelling, Technical University of Denmark, DTU, Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, Tech. Rep., may 2004.

[49] M. Weber. (1999) Caltech faces 1999. Available at: http://www.vision.caltech.edu/html-files/archive.html.

[50] M. H. H. Ara V. Nefian, Mehdi Khosravi, "Real-time detection of human faces in uncontrolled environments," pp. 3024 – 3024 – 9, 1997.

[51] C. Thomaz and G. Giraldi, "A new ranking method for principal components analysis and its application to face image analysis," *Image Vision Comput.*, vol. 28, pp. 902–913, 06 2010.

[52] P. J. Phillips, H. Moon, S. A. Rizvi, and P. J. Rauss, "The feret evaluation methodology for face-recognition algorithms," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 10, pp. 1090–1104, Oct. 2000.

[53] T. Sim, S. Baker, and M. Bsat, "The cmu pose, illumination, and expression (pie) database of human faces," Carnegie Mellon University, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-01-02, January 2001.

[54] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, "Labeled faces in the wild: A database for studying face recognition in unconstrained environments," University of Massachusetts, Amherst, Tech. Rep. 07-49, October 2007.

[55] G. B. Huang and E. Learned-Miller, "Labeled faces in the wild: Updates and new reporting procedures," University of Massachusetts, Amherst, Tech. Rep. UM-CS-2014-003, May 2014.

[56] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "Deepface: Closing the gap to human-level performance in face verification," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014.

[57] M. Osadchy, B. Pinkas, A. Jarrous, and B. Moskovich, "Scifi - a system for secure face identification," in *2010 IEEE Symposium on Security and Privacy*, May 2010, pp. 239–254.

[58] M. Turk and A. Pentland, "Eigenfaces for recognition," *J. Cognitive Neuroscience*, vol. 3, no. 1, pp. 71–86, Jan. 1991.