Hybrid Protection of Biometric Templates by Combining Homomorphic Encryption and Cancelable Biometrics

Hatef Otroshi Shahreza[1,2], Christian Rathgeb[3], Dailé Osorio-Roig[3],
Vedrana Krivokuća Hahn[1], Sébastien Marcel[1,4], and Christoph Busch[3,5]

[1]Idiap Research Institute, Switzerland
[2]École Polytechnique Fédérale de Lausanne (EPFL), Switzerland
[3]Hochschule Darmstadt, Germany
[4]Université de Lausanne (UNIL), Switzerland
[5]Norwegian University of Science and Technology (NTNU), Norway

Abstract

Homomorphic Encryption (HE) has become a well-known tool for privacy-preserving recognition in biometric systems. Despite some important advantages of HE (such as preservation of recognition accuracy), there are two main drawbacks in the application of HE to biometric recognition systems: first, the security of the system solely depends on the secrecy of the private (decryption) key; second, the computational costs of the operations on the ciphertexts are expensive. To address these challenges, in this paper we propose a hybrid scheme for the protection of biometric templates, which combines cancelable biometrics (CB) methods and HE. Applying CB prior to HE enhances both the security and privacy of the overall system, since the protected templates remain irreversible even if the secret keys are leaked (commonly referred to as the full disclosure scenario). In addition, we can reduce the dimensionality of templates using CB before applying HE, which speeds up the computation over the ciphertexts. We use BioHashing, Multi-Layer Perceptron (MLP) hashing, and Index-of-Maximum (IoM) hashing as different CB methods, and for each of these schemes, we propose a method for computing scores between hybrid-protected templates in the encrypted domain. We evaluate our proposed hybrid scheme using different state-of-the-art face recognition models (ArcFace, ElasticFace, and FaceNet) on the MOBIO and LFW datasets. The source code of our experiments is publicly available, so our work can be fully reproduced.

1. Introduction

Biometric recognition systems generally establish the identities of data subjects (i.e., users) by employing the features extracted from biometric samples (referred to as biometric templates). These features convey privacy-sensitive information about the identities of users enrolled in the biometric recognition system. Hence, if an adversary gains access to the unprotected biometric templates stored in the database of a biometric recognition system, they could obtain critical information about the enrolled individuals. For example, [26, 38] showed that an adversary can invert deep facial templates to reconstruct approximations of the underlying face images. Similarly, it has been shown that other biometric modalities can also be reconstructed from stored templates, e.g., vascular images from the vascular binary templates [23]. Recent data protection frameworks, such as the EU General Data Protection Regulation (GDPR) [32], also consider biometric data as sensitive information and impose legal obligations to protect this data.

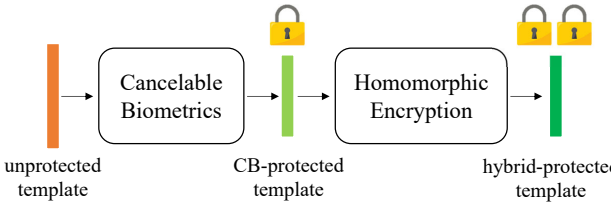

Figure 1: General scheme of the proposed hybrid protection method

To protect biometric data, several biometric template protection (BTP) schemes have been proposed in the literature [34, 33]. In general for a BTP scheme, the ISO/IEC 24745 standard [17] establishes four main requirements:

- *Renewability*: We should be able to generate a new protected template for a subject whose template is compromised.

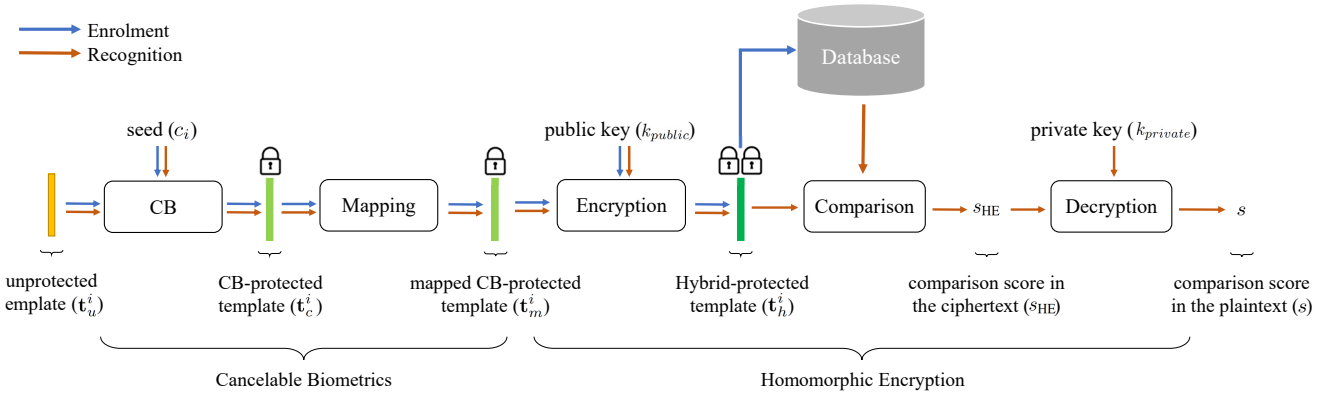

Figure 2: Block diagram of the proposed hybrid protection method

- *Unlinkability*: There should be no leakage of information between different protected templates of the same (unprotected) biometric template.

- *Irreversibility*: It should be computationally infeasible to reconstruct the original biometric templates or eventually the biometric data from the protected templates.

- *Recognition Performance*: The protected templates should be discriminative enough to be used for accurate recognition.

BTP methods are commonly categorized into *cancelable biometrics* (CB) and *biometric cryptosystems*. In cancelable biometrics protection methods (such as BioHashing [19], Multi-Layer Perceptron (MLP) Hashing [39], Index-of-Maximum (IoM) Hashing [20], etc.) a transformation function (which is dependent on a *key*) is often utilized to generate protected templates, and then the recognition is carried out by comparing the transformed templates [28, 33]. However, in biometric cryptosystems (such as fuzzy vault [21], fuzzy commitment [22], etc.), a key is either generated from a biometric template or bound with a biometric template. Then, recognition is performed based on the correct generation or retrieval of the key [42, 31].

As an alternative to cancelable biometrics and biometric cryptosystems, we could use *Homomorphic Encryption (HE)*, which allows computations to be carried out on the ciphertexts and generates encrypted results. Then, the results (which are in the encrypted domain) can be decrypted to plaintexts. The decrypted results will exactly correspond to the results of the operations performed in the plaintext domain (i.e., on the original features). Based on the allowed types and numbers of operations on the ciphertexts, the HE-based systems can also be categorized into three classes:

- *Partially Homomorphic Encryption (PHE)* systems that support only one type of arithmetic operation (i.e., either addition or multiplication) in the encrypted domain with no limit on the number of operations (e.g., [30, 11]).

- *Somewhat Homomorphic Encryption (SWHE)* systems that support both addition and multiplication but with a limited number of operations (e.g., [4, 43]).

- *Fully Homomorphic Encryption (FHE)* systems that support additions and multiplications in the encrypted domain with no limit on the number of operations (e.g., [6, 14, 13]).

Several works in the literature have used FHE for template protection in biometric recognition systems. In [3], a secure face verification system based on HE was proposed. The application of HE for face identification and face verification were also investigated in [8] and [25], respectively. In [24], HE was used for iris verification and identification. In [15], a multimodal biometric verification system using HE was proposed, and different fusion strategies were studied. There are also some works in the literature that focus on reducing computation and enhancing the efficiency of applying HE in biometric recognition systems. For example, [3] and [12] performed dimensionality reduction on the biometric features prior to HE. In [29] and [9], indexing and searching in the system's database was enhanced (for the identification application).

Despite several important advantages of HE (such as preservation of biometric recognition accuracy, as well as provable security guarantees), there are two main drawbacks in the application of HE as a BTP scheme. First, if the private (decryption) key is leaked, then the templates can be easily decrypted and inverted, which means that the security of the system solely depends on the secrecy of the keys. Second, the computational complexity of arithmetic operations on the ciphertexts is significant. To address these shortcomings, in this paper we propose a hybrid

Table 1: Protection of cancelable biometrics (CB), homomorphic encryption (HE), and hybrid (CB+HE) protection against three different threat models in the ISO/IEC 30136 standard.

Protection method	Naive/ Standard	Full disclosure
Cancelable Biometrics	✓	(✓)
Homomorphic Encryption	✓*	✗
Hybrid (CB+HE)	✓*	(✓)

*provable secure

BTP scheme using CB methods and FHE to protect biometric templates (illustrated in Fig. 1). The proposed hybrid scheme tackles the security challenge in HE-based systems when the private key is disclosed. In such cases, the CB provides more security for the system and helps to ensure that the protected templates remain irreversibile even if an attacker manages to successfully decrypt the HE-protected templates. Table 1 compares the protection of cancelable biometrics (CB), homomorphic encryption (HE), and hybrid (CB+HE) protection against three different threat models introduced in the ISO/IEC 30136 standard [18]:

- *Naive threat model* is the case where the adversary has black box knowledge about the protection method, with no further information about the underlying algorithm and any associated secrets. We can also assume that the adversary has access to a small set of protected templates (not a large biometric database).
- *Standard threat model* is the case where the adversary has full knowledge of the protection algorithm, but does not know the secrets and, therefore, cannot execute submodules that require the secrets.
- *Full disclosure threat model* refers to the case where the adversary knows everything about the system, including all the submodules and secrets.

In addition to improving the security of the protected biometric system, our experiments show that CB methods can additionally reduce the dimensionality of templates before applying HE, thereby decreasing the complexity of operations performed on the ciphertexts. The results in [37] also showed that we can reduce dimensionality of the output of BioHashing and still achieve the recognition performance of the baseline system (which uses unprotected templates). In the experiment presented in this paper, we use the following CB methods to generate protected templates prior to the application of HE: BioHashing [19], Multi-Layer Perceptron (MLP) hashing [39], and Index-of-Maximum (IoM) hashing [20]. For each of these CB schemes, we propose a method for computing scores between probe and reference templates in the encrypted domain. We evaluate our proposed hybrid scheme using different state-of-the-art (SOTA) face recognition models (i.e., ArcFace [7], ElasticFace [5], and FaceNet [35]) on the Labeled Faced in the Wild (LFW) [16] and MOBIO [27] datasets.

To elucidate the contributions of our paper, we summarize them hereunder:

- We propose a generic hybrid BTP scheme using CB methods and HE. The proposed hybrid scheme provides more security than applying HE on its own to biometric templates. In particular, in the *full disclosure* threat model [18] (where algorithms and secretes are disclosed to an adversary), the protected templates remain irreversible.
- In the proposed hybrid scheme, we can reduce the dimensionality of the biometric templates using a CB method, prior to applying HE while preserving recognition performance. This dimensianlity reduction can decrease the computations on the ciphertexts when applying HE. To the best of our knowledge, dimensionality reduction prior to HE using cancelable BTP is original and was not published before.
- We show that the scoring functions used in the comparison of templates protected via CB schemes (such as BioHashing, MLP-Hashing, and IoM Hashing) can be adapted to perform equivalent computations in the HE domain. Therefore, the hybrid method achieves the same recognition performance as the CB scheme.

The remainder of this paper is structured as follows. First, we describe our hybrid biometric template protection method in section 2. Next, in section 3, we present the experiments and discuss our results. Finally, the paper is concluded in section 4.

2. Proposed template protection method

In general, the input to the proposed template protection method can be the biometric templates extracted from different biometric modalities (e.g., face, speech, fingerprint, iris, finger vein, etc.) and with different data formats (e.g., binary, integer, float, etc.). In section 2.1, we describe the general formulation of our proposed hybrid protection method. Next, in section 2.2, we consider the combination of different CB methods (including BioHashing, MLP-Hashing, and IoM Hashing) with an HE algorithm, and we describe our hybrid template protection scheme.

2.1. General formulation

Let $\mathbf{t}_u^i$ denote the unprotected template extracted from the biometric data of the subject i. We can generate the CB-protected template $\mathbf{t}_c^i$ using the CB method $\mathcal{P}(.,.)$ applied on the unprotected template $\mathbf{t}_u^i$ along with the seed c_i:

$$\mathbf{t}_c^i = \mathcal{P}(\mathbf{t}_u^i, c_i) \tag{1}$$

To encrypt the CB-protected template $\mathbf{t}_c^i$ using HE, we may need to perform a pre-processing step prior to encoding. Therefore, we can define a mapping function $M_{\mathcal{P}}(.)$ to change the representation of the CB-protected template $\mathbf{t}_c^i$ and generate the mapped CB-protected template $\mathbf{t}_m^i$:

$$\mathbf{t}_m^i = M_{\mathcal{P}}(\mathbf{t}_c^i) \tag{2}$$

Next, we can generate the hybrid-protected template $\mathbf{t}_h^i$ (i.e., the ciphertext) by applying HE-based encryption function $\text{Enc}_{\text{HE}}(.,.)$ on the mapped CB-protected template $\mathbf{t}_m^i$, using the public key k_{public}:

$$\mathbf{t}_h^i = \text{Enc}_{\text{HE}}(\mathbf{t}_m^i, k_{public}) \tag{3}$$

In the enrolment stage, the hybrid-protected template $\mathbf{t}_h^i$ is then stored in the system database as the reference template. In the recognition stage, the hybrid-protected template of the probe should be compared to the reference templates in the homomorphically encrypted domain (i.e., the comparison should be between the corresponding ciphertexts). To calculate the comparison score between the hybrid-protected probe template $\mathbf{t}_h^{probe}$ and each hybrid-protected reference template $\mathbf{t}_h^{ref}$, we should employ an appropriate function $\text{Comp}_{\text{HE}}^{\mathcal{P}}(.,.)$, which corresponds to the utilized CB method P, in the encrypted domain. Hence, we need to compute the score between the reference and probe ciphertexts as follows:

$$s_{\text{HE}} = \text{Comp}_{\text{HE}}^{\mathcal{P}}(\mathbf{t}_h^{probe}, \mathbf{t}_h^{ref}) \tag{4}$$

Finally, we can decrypt the encrypted score s_{HE} to the plaintext using the private key $k_{private}$ as below:

$$s = \text{Dec}_{\text{HE}}(s_{\text{HE}}, k_{private}), \tag{5}$$

where $\text{Dec}_{\text{HE}}(.,.)$ denotes the decryption function of HE. Fig.2 illustrates the block diagram of the proposed hybrid BTP scheme. In the subsequent experiments, we will evaluate the proposed protection method on face recognition systems operating in verification mode only.

2.2. Combinations of different CB methods with HE

In the proposed hybrid protection scheme, we can generally use different CB methods and different HE algorithms. In this paper, we employ three different CB methods, including BioHashing [19], Multi-Layer Perceptron (MLP) Hashing [39], and Index-of-Maximum (IoM) Hashing [20]. For HE, we use the Brakerski/Fan-Vercauteren (BFV) scheme [13], which supports homomorphic operations on integer templates. Since the aforementioned CB methods generate binary and integer values, we do not need to perform quantization on the CB-protected templates (unlike when applying HE on unprotected templates that may contain floating point values). However, we might perform a mapping (i.e., $M_{\mathcal{P}}(.)$) to change the representation of the CB-protected templates prior to applying HE so that the corresponding comparison function $\text{Comp}_{\text{HE}}^{\mathcal{P}}(.,.)$ can be properly applied on the hybrid-protected templates in the encrypted domain. Hereunder, we describe the application of BioHashing, MLP-Hashing, and IoM Hashing in our proposed method:

BioHashing and MLP-Hashing BioHashing and MLP-Hashing CB methods generate binary-valued templates and use Hamming distance for calculating the comparison scores during recognition [39]. Hence, we propose to encrypt the binary-valued templates generated by these CB methods directly during the HE protection stage, with no further mapping (i.e., $\mathbf{t}_m^i = M_{\mathcal{P}}(\mathbf{t}_c^i) = \mathbf{t}_c^i$). Then, we can apply equivalent homomorphic operations to calculate the sum squared error for $\text{Comp}_{\text{HE}}^{P}(.,.)$ on the hybrid-protected templates.

IoM Hashing The IoM Hashing CB scheme generates integer-valued templates and uses the average number of collisions for calculating the comparison scores during recognition [39]. Therefore, we propose to represent each integer element of IoM-hashed templates using one-hot encoding prior to encrypting them via HE (i.e., by one-hot encoding each integer element of IoM-Hash is mapped to a vector of zeros and a single one, where the index of the single one corresponds to the value of the IoM-Hash element). Therefore, $M_{\mathcal{P}}(.)$ will be a one-hot encoding (i.e., $M_{\mathcal{P}}(t_c^i) = \text{OneHot}(t_c^i)$). Then, for comparison function $\text{Comp}_{\text{HE}}^{\mathcal{P}}(.,.)$ we can apply a series of homomorphic operations, which is equivalent to calculating the sum squared error between the probe and reference hybrid-protected templates.

3. Experiments

In this section, we describe our experiments and evaluate the proposed hybrid BTP scheme. In section 3.1, we first detail our experimental setup. In section 3.2, we analyze the recognition performance and execution time of the proposed method in different scenarios and with different configurations. Finally, we discuss our experimental results in section 3.3. We should note that in this paper we do not evaluate the renewability, unlinkablity, and irreversibility characteristics of our hybrid method, since these requirements have already been shown to be satisfied by the adopted CB methods and HE in the literature (e.g., [39, 20, 8]).

3.1. Experimental Setup

Baseline methods In our experiments, we use three SOTA face recognition models[1], including ArcFace [7], ElasticFace [5], and FaceNet [35]. As our baseline methods, we consider applying HE on the extracted embeddings (without first applying CB). Therefore, for the BFV HE algorithm, we need to quantize the embeddings prior to HE in order to obtain integer values. In our experiments, we use the equal-probable quantile quantization scheme [10] with 4 quantization levels.

Evaluation Datasets We use the MOBIO [27] and Labeled Faced in the Wild (LFW) [16] databases to evaluate the recognition performance of the proposed hybrid BTP method on SOTA face recognition models. The MOBIO dataset is a bimodal dataset consisting of face and audio data acquired using mobile devices from 150 people. In our experiments, we use the *development* subset of the *mobio-male* protocol[2]. The LFW database contains 13,233 face images of 5,749 subjects, where 1,680 subjects have two or more face images. We use the *View 2* protocol[3] in our experiments.

Evaluation Scenarios To evaluate the recognition performance of our hybrid BTP method, we consider two scenarios in our experiments: the *normal* scenario and the *full disclosure* scenario. The *normal* scenario is the expected scenario in practice, where users' keys (for the CB scheme) and HE keys are secret. On the other hand, the *full disclosure* scenario (corresponds to the *full disclosure threat model* in the ISO/IEC 30136 standard [18]) is the case where we assume that everything about the system (including the protection algorithm, as well as all submodules and secrets) is disclosed. In particular, the HE keys are leaked, and we also assume that the adversary knows the users' keys for the CB schemes.

Implementation details and source code To implement the biometric recognition pipeline in our experiments, we use the Bob[4] toolbox [2, 1]. In addition, we use the open-source implementations (in Bob) of the BioHashing and MLP-Hash CB methods [41, 40, 39]. For the IoM Hashing CB method, we adopt the open-source implementation of the Gaussian Random Projection-based Hashing (IoM-GRP Hashing) algorithm [20] in Bob [39] with 3 Gaussian random matrices. Hence, each element in IoM-protected templates belongs to $\{0, 1, 2\}$, and therefore we can use one-hot encoding of 3 bits length as our mapping function for implementing the proposed hybrid template protection method. In the following experiments, if the length of the CB-protected templates $\mathbf{t}_c^i$ is not specified, the length is equal to the length of the unprotected template $\mathbf{t}_u^i$. To implement the BFV algorithm, we use the SEAL-Python[5] wrapper on Python 3.8, which uses the C++ SEAL open-source library [36]. The execution times reported in this paper are measured on a system equipped with an Intel(R) Core(TM) i7-7700K CPU @ 4.20GHz. The source code from our experiments is publicly available to help reproduce our results[6].

3.2. Analysis

Fig. 3 compares the Receiver Operating Characteristic (ROC) curves of unprotected, HE-protected, CB-protected, and hybrid-protected (using our proposed BTP scheme) templates of ArcFace on the LFW dataset for different CB methods (i.e., BioHashing, MLP-Hashing, and IoM Hashing) in the *normal* and *full disclosure* scenarios. As this figure shows, the proposed hybrid method achieves exactly the same performance as the CB-protected templates in the *normal* and *full disclosure* scenarios for all CB methods. In addition, the hybrid-protected templates have a marginal improvement to the unprotected templates in the *normal* scenario. Comparing with HE-protected templates, in the *normal* scenario the hybrid-protected templates have slightly better performance for high values of the False Match Rate (FMR) and slightly worse performance for low values of the FMR. However, in each case the performance attainable using hybrid-protected templates is fairly close to that attainable using HE-protected templates. In the *full disclosure* scenario, while the recognition performance of HE-protected templates remains similar to the *normal* scenario, the performance of CB-protected templates degrades.

Table 2 reports the average execution time (over 100 executions) and recognition performance of HE and also the proposed hybrid method in the *normal* and *full disclosure* scenarios on the MOBIO and LFW datasets, when the adopted CB method is BioHashing and the length of the CB-protected templates (i.e., BioHashes) varies. In this table, $\ell_{\mathbf{t}_m}$ indicates the length of the mapped CB-protected template and α denotes the ratio of the length of the CB-protected template $\ell_{\mathbf{t}_c}$ to the length of the unprotected template $\ell_{\mathbf{t}_u}$ (i.e., $\alpha = \ell_{\mathbf{t}_c}/\ell_{\mathbf{t}_u}$). Tables 3 and 4 also report similar evaluation when applying MLP-Hashing and IoM Hashing, respectively, in our proposed hybrid method. As these tables show, in general, the hybrid-protected templates

[1]The implementation of each face recognition model is available at `https://gitlab.idiap.ch/bob/bob.bio.face`

[2]The implementation of the *mobio-male* protocol for the MOBIO dataset is available at `https://gitlab.idiap.ch/bob/bob.db.mobio`

[3]The implementation of the *View 2* protocol for the LFW dataset is available at `https://gitlab.idiap.ch/bob/bob.db.lfw`

[4]Available at `https://www.idiap.ch/software/bob/`

[5]Available at `https://github.com/Huelse/SEAL-Python`

[6]Source code: `https://gitlab.idiap.ch/bob/bob.paper.ijcb2022_hybrid_btp`

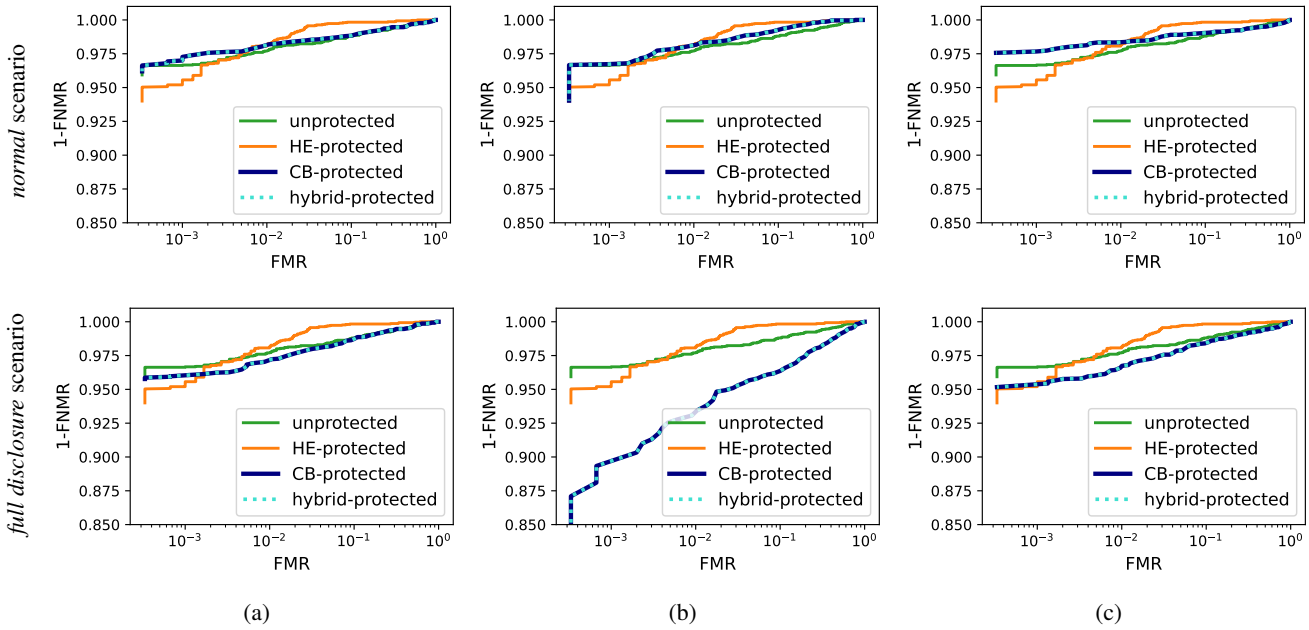

Figure 3: ROC curves of the unprotected, HE-protected, CB-protected, and hybrid-protected versions of features extracted by the ArcFace model on the LFW dataset in the (a) *normal* (first row) and *full disclosure* (second row) scenarios using (a) BioHasing, (b) MLP-Hashing, and (c) IoM Hashing.

Table 2: The average execution time (milliseconds) and recognition performance (in terms of TMR at FMR = 0.001) of HE and the proposed hybrid method, when applying **BioHashing** in the *normal* and *full disclosure* scenarios, on the MOBIO and LFW datasets using different face recognition models. In each model, the first row indicates HE protection (no CB) and the other rows show our hybrid template protection.

FR Model	α	$\ell_{\mathbf{t}_m}$	**Average Execution Time (ms)**					***normal* scenario**		***full disclosure* scenario**	
			CB	Encoding	Comparison	Decoding	Total	MOBIO	LFW	MOBIO	LFW
ArcFace ($\ell_{\mathbf{t}_u}$=512)	-	-	-	1.24 ± 0.55	329.04 ± 3.96	0.38 ± 0.00	330.66 ± 4.50	100.00%	95.20%	100.00%	95.20%
	1.00	512	15.08 ± 2.42	1.24 ± 0.55	329.04 ± 3.96	0.38 ± 0.00	345.74 ± 5.11	100.00%	97.00%	100.00%	95.93%
	0.75	384	9.47 ± 1.26	1.28 ± 0.04	263.67 ± 1.71	0.42 ± 0.01	274.82 ± 2.13	100.00%	96.20%	100.00%	95.13%
	0.50	256	6.24 ± 1.92	1.19 ± 0.04	166.21 ± 4.40	0.38 ± 0.00	174.01 ± 4.83	100.00%	94.83%	100.00%	93.53%
	0.25	128	1.77 ± 0.76	1.19 ± 0.01	84.43 ± 0.79	0.38 ± 0.00	87.75 ± 1.10	100.00%	84.83%	99.68%	84.67%
ElasticFace ($\ell_{\mathbf{t}_u}$=512)	-	-	-	1.24 ± 0.55	329.04 ± 3.96	0.38 ± 0.00	330.66 ± 4.50	99.96%	87.97%	99.96%	87.97%
	1.00	512	15.08 ± 2.42	1.24 ± 0.55	329.04 ± 3.96	0.38 ± 0.00	345.74 ± 5.11	100.00%	96.47%	100.00%	93.43%
	0.75	384	9.47 ± 1.26	1.28 ± 0.04	263.67 ± 1.71	0.42 ± 0.01	274.82 ± 2.13	100.00%	95.77%	100.00%	95.13%
	0.50	256	6.24 ± 1.92	1.19 ± 0.04	166.21 ± 4.40	0.38 ± 0.00	174.01 ± 4.83	100.00%	94.63%	99.88%	88.43%
	0.25	128	1.77 ± 0.76	1.19 ± 0.01	84.43 ± 0.79	0.38 ± 0.00	87.75 ± 1.10	99.68%	86.93%	86.27%	80.67%
FaceNet ($\ell_{\mathbf{t}_u}$=128)	-	-	-	1.19 ± 0.01	84.27 ± 0.08	0.38 ± 0.00	85.83 ± 0.08	98.41%	88.97%	98.41%	88.97%
	1.00	128	1.2 ± 0.69	1.19 ± 0.01	84.27 ± 0.08	0.38 ± 0.00	87.03 ± 0.70	99.92%	95.33%	92.86%	78.33%
	0.75	96	0.49 ± 0.65	1.19 ± 0.01	64.22 ± 2.91	0.38 ± 0.00	66.28 ± 2.98	99.64%	88.73%	73.93%	75.60%
	0.50	64	0.36 ± 0.17	1.19 ± 0.06	43.71 ± 1.11	0.38 ± 0.02	45.64 ± 1.13	98.69%	73.87%	83.69%	52.47%
	0.25	32	0.12 ± 0.01	1.19 ± 0.01	23.14 ± 0.05	0.38 ± 0.00	24.82 ± 0.05	85.83%	48.3%	32.26%	23.17%

can achieve superior recognition performance compared to the HE-protected templates in the *normal* scenario. In the *full disclosure* scenario, hybrid-protected templates (with $\alpha = 1$) have competitive performance with HE-protected templates. Notwithstanding of the good performance of HE-protected templates in the *full disclosure* scenario, we

Table 3: The average execution time (milliseconds) and recognition performance (in terms of TMR at FMR = 0.001) of HE and the proposed hybrid method, when applying **MLP-Hashing** in the *normal* and *full disclosure* scenarios, on the MOBIO and LFW datasets using different face recognition models. In each model, the first row indicates HE protection (no CB) and the other rows show our hybrid template protection.

FR Model	α	$\ell_{\mathbf{t}_m}$	**Average Execution Time (ms)**					***normal* scenario**		***full disclosure* scenario**	
			CB	Encoding	Comparison	Decoding	Total	MOBIO	LFW	MOBIO	LFW
ArcFace ($\ell_{\mathbf{t}_u}$=512)	-	-	-	1.24 ± 0.55	329.04 ± 3.96	0.38 ± 0.00	330.66 ± 4.50	100.00%	95.20%	100.00%	95.20%
	1.00	512	56.07 ± 10.83	1.19 ± 0.04	328.30 ± 0.32	0.39 ± 0.00	385.93 ± 10.83	100.00%	96.73%	99.84%	88.10%
	0.75	384	48.39 ± 10.25	1.19 ± 0.04	247.20 ± 1.87	0.39 ± 0.01	297.16 ± 10.42	100.00%	96.43%	99.76%	88.33%
	0.50	256	39.86 ± 5.09	1.24 ± 0.02	172.30 ± 0.76	0.41 ± 0.01	213.80 ± 5.15	100.00%	93.27%	98.61%	80.87%
	0.25	128	37.59 ± 6.23	1.23 ± 0.03	87.03 ± 1.56	0.40 ± 0.01	126.24 ± 6.43	98.77%	86.57%	85.40%	53.37%
ElasticFace ($\ell_{\mathbf{t}_u}$=512)	-	-	-	1.24 ± 0.55	329.04 ± 3.96	0.38 ± 0.00	330.66 ± 4.50	99.96%	87.97%	99.96%	87.97%
	1.00	512	56.07 ± 10.83	1.19 ± 0.04	328.30 ± 0.32	0.39 ± 0.00	385.93 ± 10.83	100.00%	94.50%	99.68%	87.97%
	0.75	384	48.39 ± 10.25	1.19 ± 0.04	247.20 ± 1.87	0.39 ± 0.01	297.16 ± 10.42	100.00%	95.07%	99.60%	82.43%
	0.50	256	39.86 ± 5.09	1.24 ± 0.02	172.30 ± 0.76	0.41 ± 0.01	213.80 ± 5.15	100.00%	92.00%	97.90%	69.17%
	0.25	128	37.59 ± 6.23	1.23 ± 0.03	87.03 ± 1.56	0.40 ± 0.01	126.24 ± 6.43	99.01%	78.17%	74.56%	44.37%
FaceNet ($\ell_{\mathbf{t}_u}$=128)	-	-	-	1.19 ± 0.01	84.27 ± 0.08	0.38 ± 0.00	85.83 ± 0.08	98.41%	88.97%	98.41%	88.97%
	1.00	128	1.62 ± 0.24	1.17 ± 0.01	84.11 ± 0.08	0.37 ± 0.01	87.27 ± 0.25	99.33%	86.53%	47.34%	50.53%
	0.75	96	1.51 ± 0.27	1.17 ± 0.01	63.76 ± 0.06	0.36 ± 0.00	66.8 ± 0.28	98.57%	79.90%	47.54%	35.07%
	0.50	64	1.43 ± 0.29	1.17 ± 0.01	43.44 ± 0.04	0.36 ± 0.00	46.40 ± 0.29	92.34%	55.73%	49.33%	25.50%
	0.25	32	1.31 ± 0.29	1.17 ± 0.01	23.06 ± 0.03	0.36 ± 0.00	25.90 ± 0.29	58.69%	31.43%	20.99%	9.90%

Table 4: The average execution time (milliseconds) and recognition performance (in terms of TMR at FMR = 0.001) of HE and the proposed hybrid method, when applying **IoM Hashing** in the *normal* and *full disclosure* scenarios, on the MOBIO and LFW datasets using different face recognition models. In each model, the first row indicates HE protection (no CB) and the other rows show our hybrid template protection.

FR Model	α	$\ell_{\mathbf{t}_m}$	**Average Execution Time (ms)**					***normal* scenario**		***full disclosure* scenario**	
			CB	Encoding	Comparison	Decoding	Totall	MOBIO	LFW	MOBIO	LFW
ArcFace ($\ell_{\mathbf{t}_u}$=512)	-	-	-	1.24 ± 0.55	329.04 ± 3.96	0.38 ± 0.00	330.66 ± 4.50	100.00%	95.20%	100.00%	95.20%
	1.00	1536	26.89 ± 2.54	1.19 ± 0.03	981.46 ± 3.23	0.39 ± 0.00	1009.92 ± 4.12	100.00%	97.67%	99.76%	95.30%
	0.75	1152	23.06 ± 8.31	1.19 ± 0.01	736.24 ± 0.82	0.38 ± 0.00	760.86 ± 8.36	100.00%	97.17%	99.76%	94.17%
	0.50	768	14.92 ± 2.07	1.19 ± 0.01	491.80 ± 0.41	0.38 ± 0.00	508.28 ± 2.11	100.00%	95.73%	99.76%	94.17%
	0.25	384	6.67 ± 0.50	1.19 ± 0.01	248.39 ± 9.48	0.38 ± 0.00	256.62 ± 9.49	100.00%	91.33%	98.93%	90.37%
ElasticFace ($\ell_{\mathbf{t}_u}$=512)	-	-	-	1.24 ± 0.55	329.04 ± 3.96	0.38 ± 0.00	330.66 ± 4.50	99.96%	87.97%	99.96%	87.97%
	1.00	1536	26.89 ± 2.54	1.19 ± 0.03	981.46 ± 3.23	0.39 ± 0.00	1009.92 ± 4.12	100.00%	96.83%	98.10%	92.63%
	0.75	1152	23.06 ± 8.31	1.19 ± 0.01	736.24 ± 0.82	0.38 ± 0.00	760.86 ± 8.36	100.00%	95.43%	98.10%	92.30%
	0.50	768	14.92 ± 2.07	1.19 ± 0.01	491.80 ± 0.41	0.38 ± 0.00	508.28 ± 2.11	100.00%	94.07%	98.10%	91.23%
	0.25	384	6.67 ± 0.50	1.19 ± 0.01	248.39 ± 9.48	0.38 ± 0.00	256.62 ± 9.49	100.00%	91.53%	98.21%	81.90%
FaceNet ($\ell_{\mathbf{t}_u}$=128)	-	-	-	1.19 ± 0.01	84.27 ± 0.08	0.38 ± 0.00	85.83 ± 0.08	98.41%	88.97%	98.41%	88.97%
	1.00	384	1.51 ± 0.02	1.19 ± 0.01	247.58 ± 1.71	0.39 ± 0.00	250.65 ± 1.71	99.96%	97.20%	95.44%	77.83%
	0.75	288	1.13 ± 0.01	1.19 ± 0.01	186.34 ± 0.63	0.38 ± 0.00	189.03 ± 0.63	99.84%	95.37%	93.61%	74.10%
	0.50	192	0.75 ± 0.01	1.19 ± 0.01	125.09 ± 0.12	0.38 ± 0.00	127.40 ± 0.12	99.33%	88.67%	87.38%	60.73%
	0.25	96	0.38 ± 0.00	1.19 ± 0.01	63.91 ± 0.10	0.38 ± 0.00	65.85 ± 0.10	91.39%	67.97%	56.39%	45.00%

should note that in this scenario the adversary can easily reconstruct the unprotected templates using the HE private (decryption) key (i.e, very poor protection). However, for hybrid-protected templates, the adversary can only reconstruct the (mapped) CB-protected templates using the HE private key, but it is still difficult for the adversary to reconstruct the unprotected templates from the CB-protected templates. We can also see that with $\alpha = 1$, the hybrid protection requires a longer execution time than HE. However, we can adjust the value of α so that the hybrid protection

achieves a shorter execution time with comparable recognition performance.

3.3. Discussion

Our experiments in section 3.2 show that the proposed hybrid scheme achieves exactly the same recognition performance as the corresponding CB method. Our experiments also show that in the *normal* scenario, the proposed hybrid method (with $\alpha = 1$) achieves superior performance compared to HE. In the *full disclosure* scenario, hybrid-protected templates (with $\alpha = 1$) achieve comparable performance with HE-protected templates for ArcFace and ElasticFace, but HE-protected templates perform better than hybrid-protected templates for FaceNet. Having said that, it is important to keep in mind that HE-protected templates can be easily inverted to recover the original (unprotected) templates, whereas hybrid-protected templates are not easily invertible due to the extra layer of protection provided by the CB method that is applied prior to HE.

Tables 2-4 also show that there is a trade-off between the execution time and recognition performance when using the proposed hybrid protection method. This trade-off can be controlled with α. For $\alpha = 1$, hybrid-protected templates require longer execution times than HE-protected templates. However, with smaller α, CB can in practice reduce the dimensionality of features prior to HE. Therefore, we can achieve a shorter execution time compared to HE. In particular, for a proper choice of α, for the hybrid-protected templates we can simultaneously achieve a shorter execution time and comparable performance to the HE-protected templates. For example, for ElasticFace and BioHashing, we could set $\alpha = 0.25$, whereas for FaceNet at the same setting hybrid-protected templates have worse performance than HE-protected templates. Therefore, in this case, it would be better to set α to a higher value such as $\alpha = 0.75$. The suitable lengths of BioHash-protected templates (and therefore α), which maintain the recognition performance of unprotected templates, are investigated in [37] for different SOTA face recognition models.

4. Conclusion

In this paper, we proposed a generic hybrid BTP scheme for biometric templates by combining Cancelable Biometrics (CB) and Homomorphic Encryption (HE). We showed that the comparison methods of CB schemes (such as BioHashing, MLP-Hashing, and IoM Hashing) can be adapted to perform equivalent computations in the HE domain, and therefore our hybrid scheme was found to achieve equal recognition performance with the corresponding CB. Our experiments further showed that the proposed hybrid method is able to achieve better performance compared to HE alone in the *normal* scenario. In the *full disclosure* threat model (where algorithms and secretes are disclosed to an adversary), the hybrid-protected templates were found to have comparable performance with HE-protected templates in most cases, when the length of the CB-protected templates was equal to the length of the unprotected templates. As the length of the CB-protected templates was decreased, the performance of the hybrid-protected templates was found to also decrease, so for much smaller lengths the performance of templates protected using HE alone was sometimes found to be better than that of hybrid-protected templates. However, the main drawback of HE-protected templates is that they can be easily inverted by an adversary with access to the secret decryption key, while hybrid-protected templates remain irreversible in this case. Besides the additional template protection offered by our hybrid BTP method, it is also useful for reducing the dimensionality of the biometric templates with CB, prior to applying HE, which can decrease the amount of computation on the encrypted templates (ciphertexts). In particular, by appropriately tuning the length of CB-protected templates, we could achieve comparable recognition performance with HE, but with a faster execution time.

Acknowledgment

This research is based upon work supported by the H2020 TReSPAsS-ETN Marie Skłodowska-Curie early training network (grant agreement 860813) and by the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE. The authors would like to thank Tiago de Freitas Pereira (Idiap Research Institute, Switzerland) for his support in the code implementations using the Bob toolbox.

References

[1] A. Anjos, M. Günther, T. de Freitas Pereira, P. Korshunov, A. Mohammadi, and S. Marcel. Continuously reproducing toolchains in pattern recognition and machine learning experiments. In *Proceedings of the International Conference on Machine Learning (ICML)*, Aug. 2017.

[2] A. Anjos, L. E. Shafey, R. Wallace, M. Günther, C. McCool, and S. Marcel. Bob: a free signal processing and machine learning toolbox for researchers. In *Proceedings of the 20th ACM Conference on Multimedia Systems (ACMMM)*, Oct. 2012.

[3] V. N. Boddeti. Secure face matching using fully homomorphic encryption. In *Proceedings of the 9th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pages 1–10. IEEE, 2018.

[4] D. Boneh, E.-J. Goh, and K. Nissim. Evaluating 2-dnf formulas on ciphertexts. In *Proceedings of the Theory of Cryptography Conference*, pages 325–341. Springer, 2005.

[5] F. Boutros, N. Damer, F. Kirchbuchner, and A. Kuijper. Elasticface: Elastic margin loss for deep face recognition. pages 1578–1587, 2022.

[6] Z. Brakerski, C. Gentry, and V. Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. *ACM Transactions on Computation Theory (TOCT)*, 6(3):1–36, 2014.

[7] J. Deng, J. Guo, X. Niannan, and S. Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[8] P. Drozdowski, N. Buchmann, C. Rathgeb, M. Margraf, and C. Busch. On the application of homomorphic encryption to face identification. In *Proceedings of the International Conference of the Biometrics Special Interest Group (BIOSIG)*, pages 1–5. IEEE, 2019.

[9] P. Drozdowski, F. Stockhardt, C. Rathgeb, D. Osorio-Roig, and C. Busch. Feature fusion methods for indexing and retrieval of biometric data: Application to face recognition with privacy protection. *IEEE Access*, 9:139361–139378, 2021.

[10] P. Drozdowski, F. Struck, C. Rathgeb, and C. Busch. Benchmarking binarisation schemes for deep face templates. In *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, pages 191–195. IEEE, 2018.

[11] T. ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31(4):469–472, 1985.

[12] J. J. Engelsma, A. K. Jain, and V. N. Boddeti. Hers: Homomorphically encrypted representation search. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 2022.

[13] J. Fan and F. Vercauteren. Somewhat practical fully homomorphic encryption. *Cryptology ePrint Archive*, 2012.

[14] C. Gentry. *A fully homomorphic encryption scheme*. Stanford university, 2009.

[15] M. Gomez-Barrero, E. Maiorana, J. Galbally, P. Campisi, and J. Fierrez. Multi-biometric template protection based on homomorphic encryption. *Pattern Recognition*, 67:149–163, 2017.

[16] 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. Technical Report 07-49, University of Massachusetts, Amherst, October 2007.

[17] ISO/IEC 24745:2022(E) Information technologyy, cybersecurity and privacy protection – Biometric information protection, Feb. 2022.

[18] ISO/IEC 30136:2018(E) Information technology – Security techniques – Performance testing of biometric template protection schemes, June 2018.

[19] A. T. B. Jin, D. N. C. Ling, and A. Goh. Biohashing: two factor authentication featuring fingerprint data and tokenised random number. *Pattern Recognition*, 37(11):2245–2255, 2004.

[20] Z. Jin, J. Y. Hwang, Y.-L. Lai, S. Kim, and A. B. J. Teoh. Ranking-based locality sensitive hashing-enabled cancelable biometrics: Index-of-max hashing. *IEEE Transactions on Information Forensics and Security*, 13(2):393–407, 2017.

[21] A. Juels and M. Sudan. A fuzzy vault scheme. *Designs, Codes and Cryptography*, 38(2):237–257, 2006.

[22] A. Juels and M. Wattenberg. A fuzzy commitment scheme. In *Proceedings of the 6th ACM Conference on Computer and Communications Security*, pages 28–36, 1999.

[23] C. Kauba, S. Kirchgasser, V. Mirjalili, A. Uhl, and A. Ross. Inverse biometrics: Generating vascular images from binary templates. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 3(4):464–478, 2021.

[24] J. Kolberg, P. Bauspieß, M. Gomez-Barrero, C. Rathgeb, M. Dürmuth, and C. Busch. Template protection based on homomorphic encryption: Computationally efficient application to iris-biometric verification and identification. In *Proceedings of the IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–6. IEEE, 2019.

[25] J. Kolberg, P. Drozdowski, M. Gomez-Barrero, C. Rathgeb, and C. Busch. Efficiency analysis of post-quantum-secure face template protection schemes based on homomorphic encryption. In *Proceedings of the International Conference of the Biometrics Special Interest Group (BIOSIG)*, pages 1–4. IEEE, 2020.

[26] G. Mai, K. Cao, P. C. Yuen, and A. K. Jain. On the reconstruction of face images from deep face templates. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(5):1188–1202, 2018.

[27] C. McCool, R. Wallace, M. McLaren, L. El Shafey, and S. Marcel. Session variability modelling for face authentication. *IET Biometrics*, 2(3):117–129, Sept. 2013.

[28] K. Nandakumar and A. K. Jain. Biometric template protection: Bridging the performance gap between theory and practice. *IEEE Signal Processing Magazine*, 32(5):88–100, 2015.

[29] D. Osorio-Roig, C. Rathgeb, P. Drozdowski, and C. Busch. Stable hash generation for efficient privacy-preserving face identification. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 2021.

[30] P. Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques*, pages 223–238. Springer, 1999.

[31] C. Rathgeb, J. Merkle, J. Scholz, B. Tams, and V. Nesterowicz. Deep face fuzzy vault: Implementation and performance. *Computers & Security*, 113:102539, 2022.

[32] G. D. P. Regulation. Regulation EU 2016/679 of the european parliament and of the council of 27 april 2016. *Official Journal of the European Union*, 2016.

[33] M. Sandhya and M. V. Prasad. Biometric template protection: A systematic literature review of approaches and modalities. In *Biometric Security and Privacy*, pages 323–370. Springer, 2017.

[34] A. Sarkar and B. K. Singh. A review on performance, security and various biometric template protection schemes for biometric authentication systems. *Multimedia Tools and Applications*, pages 1–56, 2020.

[35] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Pro-*

ceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 815–823, 2015.

[36] Microsoft SEAL (release 3.6). `https://github.com/Microsoft/SEAL`, Nov. 2020. Microsoft Research, Redmond, WA.

[37] H. O. Shahreza, V. K. Hahn, and S. Marcel. On the recognition performance of biohashing on state-of-the-art face recognition models. In *Proceedings of the IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–6. IEEE, 2021.

[38] H. O. Shahreza, V. K. Hahn, and S. Marcel. Face reconstruction from deep facial embeddings using a convolutional neural network. In *Proceedings of the IEEE International Conference on Image Processing (ICIP)*. IEEE, 2022.

[39] H. O. Shahreza, V. K. Hahn, and S. Marcel. Mlp-hash: Protecting face templates via hashing of randomized multi-layer perceptron. *arXiv preprint arXiv:2204.11054*, 2022.

[40] H. O. Shahreza and S. Marcel. Deep auto-encoding and biohashing for secure finger vein recognition. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2585–2589. IEEE, 2021.

[41] H. O. Shahreza and S. Marcel. Towards protecting and enhancing vascular biometric recognition methods via biohashing and deep neural networks. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 3(3):394–404, 2021.

[42] U. Uludag, S. Pankanti, S. Prabhakar, and A. K. Jain. Biometric cryptosystems: issues and challenges. *Proceedings of the IEEE*, 92(6):948–960, 2004.

[43] A. C. Yao. Protocols for secure computations. In *Proceedings of the 23rd Annual Symposium on Foundations of Computer Science (SFCS)*, pages 160–164. IEEE, 1982.