MLP-Hash: Protecting Face Templates via Hashing of Randomized Multi-Layer Perceptron

Hatef Otroshi Shahreza[1,2], Vedrana Krivokuća Hahn[1], and Sébastien Marcel[1,3]
[1]Idiap Research Institute, Switzerland
[2]École Polytechnique Fédérale de Lausanne (EPFL), Switzerland
[3]Université de Lausanne (UNIL), Switzerland

***Abstract*—Applications of face recognition systems for authentication purposes are growing rapidly. Although state-of-the-art (SOTA) face recognition systems have high recognition accuracy, the features which are extracted for each user and are stored in the system's database contain privacy-sensitive information. Accordingly, compromising this data would jeopardize users' privacy. In this paper, we propose a new cancelable template protection method, dubbed MLP-hash, which generates protected templates by passing the extracted features through a user-specific randomly-weighted multi-layer perceptron (MLP) and binarizing the MLP output. We evaluated the unlinkability, irreversibility, and recognition accuracy of our proposed biometric template protection method to fulfill the ISO/IEC 30136 standard requirements. Our experiments with SOTA face recognition systems on the MOBIO and LFW datasets show that our method has competitive performance with the BioHashing and IoM Hashing (IoM-GRP and IoM-URP) template protection algorithms. We provide an open-source implementation of all the experiments presented in this paper so that other researchers can verify our findings and build upon our work.**

***Index Terms*—Biometrics, Face recognition, Hashing, Multi-Layer Perceptron (MLP), Template Protection.**

I. Introduction

Face recognition has become a popular authentication tool and has been widely used in recent years. The state-of-the-art (SOTA) face recognition systems mainly use convolutional neural networks (CNNs) to extract features, called "embeddings", from face images. In the enrollment stage, these features are extracted from each user's face and are stored as reference templates in the database of the face recognition system. Then, in the recognition stage, similar features are extracted from the user, and the resulting probe template is compared with the reference embedding stored in the system's database. These face features contain privacy-sensitive information about the user's identity [1], [2]. Hence, data protection regulations, such as the EU General Data Protection Regulation (GDPR) [3], consider biometric templates as sensitive data that must (legally) be protected.

To protect biometric templates, different methods have been proposed in the literature [4], [5]. According to the ISO/IEC 30136 standard [6], each biometric template protection (BTP) scheme generally should have four main properties:

This research is based upon work supported by the H2020 TReSPAsS-ETN Marie Skłodowska-Curie early training network (grant agreement 860813).

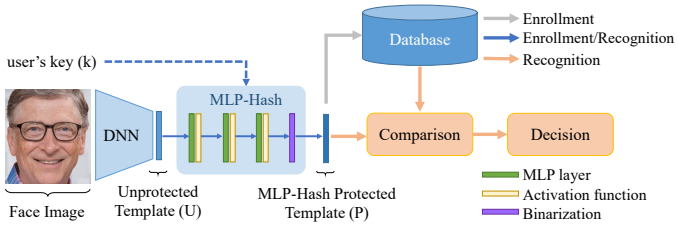

Fig. 1: Block diagram of MLP-Hash protected face recognition system

- **Cancelability**: If a biometric template is compromised, we should be able to cancel the enrolled protected template and replace it with a new protected template.
- **Unlinkability**: Considering the cancelability property, there should be no link between different protected templates from the same unprotected (original) biometric template.
- **Irreversibility**: It should be computationally difficult or impossible to recover the original biometric templates from the protected templates.
- **Recognition Accuracy**: The protected templates should allow for accurate recognition and should not result in recognition accuracy degradation.

BTP methods can generally be categorized into *cancelable biometrics* and *biometric cryptosystems*. In cancelable template protection methods (such as BioHashing [7], Index-of-Maximum (IoM) Hashing [8], etc.) a transformation function is often used (which is dependent on a *key*) to generate protected templates, and then for recognition the comparison is performed in the transformed domain [4], [9]. However, in biometric cryptosystems (such as fuzzy commitment [10], fuzzy vault [11], etc.), a key is either bound with a biometric template (called key binding schemes) or generated from a biometric template (called key generation schemes). Then, recognition is based upon correct retrieval or generation of the key [12].

In this paper, we propose a new cancelable biometric template protection scheme, dubbed MLP-Hash, which includes a non-linear projection step through a user-specific randomly-weighted multi-layer perceptron (MLP), followed by a binarization step. We employ the user's private key to initialize the MLP with random orthonormal values. Then, we project the templates to a new space through the initialized

MLP, which contains nonlinear activation functions. Finally, at the output layer, we binarize the final layer of the MLP to generate the protected template.

We evaluate the unlinkability and irreversibility properties of our template protection method to fulfill the ISO/IEC 30136 standard [6] requirements. We also consider two scenarios when evaluating the method's recognition accuracy: the *normal* scenario (which is the expected scenario in practice) and the *stolen token* scenario (which is the case when the user's MLP-Hash key is stolen). Then, we evaluate the protected templates of three SOTA face recognition methods (i.e., ArcFace [13], FaceNet [14], and InceptionResnetV2-CenterLoss [15]) on the Labeled Faced in the Wild (LFW) [16] and MOBIO [17] datasets. Our experiments show that MLP-Hash achieves promising performance in protecting SOTA face recognition systems.

The rest of this paper is organized as follows. First, we describe our biometric template protection method in section II. Then, in section III, we evaluate MLP-Hash in terms of unlinkability, irreversibility, and recognition accuracy. Finally, the paper is concluded in section IV.

II. Proposed method

Figure 1 represents the block diagram of an MLP-Hash protected face recognition system. As depicted in this figure, MLP-Hash uses unprotected features, which are extracted from the user's face image, along with the user's key, to generate the protected template. In section II-A, we describe the MLP-Hash algorithm in detail. During the enrollment stage, the protected templates are stored in the system's database and are later compared with the probe template during the recognition stage as described in section II-B. We should note that compared to BTP schemes which use neural networks and require training, e.g. [18]–[20], our proposed method does not require training and the weights are specified using the user's key (as described in section II-A).

A. MLP-Hash Algorithm

Let U indicate an unprotected biometric template (i.e., embedding) extracted by a face recognition model. The MLP-Hash protected template, P, can be generated by algorithm 1 using the user's key, k, and the unprotected template, U, in two steps. First, U is fed into an MLP with H hidden layers, activation function $F(.)$[1], and the pseudo-random orthonormal weights initialized with seed k. To generate pseudo-random orthonormal matrix $\mathbf{M}_{\perp\ell}$ in layer ℓ of the MLP, we first generate a pseudo-random matrix $\mathbf{M}_\ell$, and then apply the Gram-Schmidt orthonormalization process on the rows of $\mathbf{M}_\ell$. After feeding the U into the MLP with the pseudo-random orthonormal weights, in the second step, we binarize the output of MLP to generate the protected template, P.

[1]In this paper, we use the Rectified Linear Unit (ReLU) activation function which is a non-linear and many-to-one function.

Algorithm 1 MLP-Hash algorithm

Inputs:
- U : unprotected biometric template (i.e., embedding)
- H : number of MLP hidden layers
- L_{MLP} : set of lengths of MLP layers ($L_{\text{MLP}}^{(\ell)}$), including input layer ($\ell = 0$), hidden layers ($1 \leq \ell \leq H$), and output layer ($\ell = H + 1$)
- $F(.)$: activation function
- k : user's key

Output:
- $P = \{p_i | i = 1, 2, ..., L_{\text{MLP}}^{(H+1)}\}$ binary MLP-Hash protected template

Procedure:

Step 1: Passing through pseudo-random MLP

Set initial value of Γ with U

for ℓ **in** $\{1, ..., H + 1\}$ **do**

Generate a pseudo-random matrix $\mathbf{M}_\ell$ based on the user's seed (k): $\mathbf{M}_\ell \in \mathbb{R}^{L_{\text{MLP}}^{(\ell-1)} \times L_{\text{MLP}}^{(\ell)}}$.

Apply the Gram-Schmidt process on the rows of the generated pseudo-random matrix $\mathbf{M}_\ell$ to transform it into an orthonormal matrix $\mathbf{M}_{\perp\ell}$

Update value of Γ with matrix product of Γ and $\mathbf{M}_{\perp\ell}$

Update value of Γ by applying activation function $F(\Gamma)$

end for

Step 2: Binarizing the output of MLP

Compute $L_{\text{MLP}}^{(H+1)}$ bits MLP-Hash $\{p_i | i = 1, 2, ..., L_{\text{MLP}}^{(H+1)}\}$ from

$$p_i = \begin{cases} 0 & \text{if} \quad \Gamma_i \leq \tau \\ 1 & \text{if} \quad \Gamma_i > \tau \end{cases}, \qquad i = 1, ..., L_{\text{MLP}}^{(H+1)},$$

where τ is the average of Γ elements.

End Procedure

B. Comparing MLP-Hash Templates

In the enrollment stage, the reference MLP-Hash templates, P, should be stored in the system database (ideally separately). In the recognition stage, we use *Hamming* distance to calculate the score between each pair of *probe* and *reference* MLP-Hashed templates. In the subsequent experiments, we consider the MLP-Hash protected face recognition systems operating in verification mode only.

III. Experiments

In this section, we describe our experiments and evaluate the properties of MLP-Hash as a biometric template protection scheme in accordance with the ISO/IEC 30136 standard. First, in section III-A, we describe our experimental setup and the baselines used. Next, we evaluate the unlinkability, irreversibility and recognition accuracy of MLP-Hash in sections III-B, III-C, and III-D, respectively. We should note that cancelability is inherently satisfied in the MLP-Hash algorithm, since like other *cancelable* BTP methods, we can easily revoke the compromised template in the database, assign a new key for the user, and register the user with a new protected template. Finally, we discuss our experiments in section III-E.

A. Experimental Setup and Baselines

As stated in section I, in our experiments we used the MOBIO [17] and Labeled Faced in the Wild (LFW) [16]

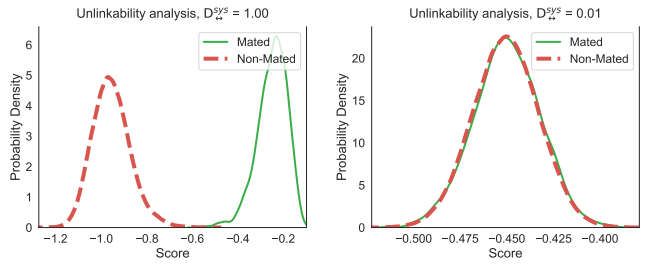

(a) (b)

Fig. 2: Unlinkability evaluation of unprotected and MLP-Hash protected ArcFace templates on the MOBIO dataset: a)Unprotected templates, b)MLP-Hash protected templates.

databases to evaluate the recognition accuracy of MLP-Hash on SOTA face recognition models. The MOBIO dataset is a bimodal dataset including audio and face data acquired using mobile devices from 152 people. We used the *development* subset of the *mobio-all* protocol[2] in our experiments. The LFW database includes 13,233 images of 5,749 people, where 1,680 people have two or more images. We used the *View 2* protocol[3] to evaluate the models. We also used three SOTA face recognition models[4], including ArcFace [13], FaceNet [14], and InceptionResnetV2-CenterLoss [15]. We compare the performance of our template protection method on the same face recognition systems with the BioHashing [7] method and two methods based on Index-of-Maximum (IoM) Hashing [8] (i.e., Gaussian random projection-based hashing, shortly GRP, and uniformly random permutation-based hashing, shortly URP). In each case, we generate protected templates whose length is equal to the length of the embedding (i.e., number of elements in the embedding) for each face recognition model. We set all the hidden layers of MLP-Hash to twice the length of the embeddings for each face recognition model. The number of hidden layers (H) was 3 in our experiments.

For our experiments, we used the Bob[5] toolbox [21], [22]. To implement the BioHashing algorithm, we used the open-source implementation of the BioHashing in Bob [23], [24]. The source code from our experiments is publicly available to help reproduce our results[6].

B. Unlinkability Evaluation

To evaluate the unlinkability criterion, we used the framework proposed in [25]. This framework uses the score distributions of the *mated* templates (i.e., different templates from the same user) and *non-mated* templates (i.e., templates from different users) to measure unlinkability with respect to the overlap of these two distributions. More particularly, with this evaluation, we expect that in the case of linkable templates, the *mated* and *non-mated* templates score distributions will be separated. However, in the case of unlinkable templates, these distributions should completely overlap. Figure 2 compares the unlinkability of original (unprotected) and MLP-Hash protected ArcFace templates on the MOBIO dataset using this evaluation framework[7]. To calculate the distribution of *mated* scores in this figure, we generated different templates for the same user using different keys, then calculated the scores between these templates. However, for the distribution of *non-mated* scores, we generated protected templates for different users (with different keys) and computed the scores between them. As shown in this figure, while the distributions of *mated* scores and *non-mated* scores are fully separated for unprotected templates, they almost completely overlap for the MLP-Hash protected templates. Furthermore, the value of the system's global unlinkability measure ($D^{sys}_{\leftrightarrow}$) is reduced from 1.0 (for the unprotected system) to 0.01 (for the MLP-Hash protected system) by deploying our template protection method, showing that the resulting protected templates are almost fully unlinkable. Table I compares the unlinkability of MLP-Hash, BioHash, IoM-GRP, and IoM-URP protected templates of the ArcFace embeddings on the MOBIO database. As this table shows, all these template protection schemes have comparable unlinkability and they are almost fully unlinkable.

TABLE I: Unlinkability evaluation of MLP-Hash, BioHash, IoM-GRP, and IoM-URP protected templates of the ArcFace embeddings in terms of the system's global unlinkability measure ($D^{sys}_{\leftrightarrow}$).

MLP-Hash	BioHash	IoM-GRP	IoM-URP
0.010	0.009	0.011	0.007

TABLE II: Irreversibility evaluation of MLP-Hash, BioHash, IoM-GRP, and IoM-URP protected templates of the ArcFace embeddings in terms of Success Attack Rate (%) on the MOBIO dataset at FMR of 10^{-2} and 10^{-3}.

Configuration	MLP-Hash	BioHash	IoM-GRP	IoM-URP
FMR = 10^{-2}	39.05	43.81	35.71	14.29
FMR = 10^{-3}	9.05	10.48	7.14	1.43

C. Irreversibility Evaluation

To evaluate the irreversibility of the proposed template protection scheme, we consider the worst-case and most difficult threat model in ISO/IEC 30136 standard (referred to as *full disclosure threat model*), where the attacker knows everything about the system, including algorithms, secret keys, etc. We assume that the attacker would invert the protected template, then use the inverted template to enter a similar unprotected system. Accordingly, we evaluate the irreversibilty in term of Success Attack Rate (SAR), which indicates the attacker's success rate in entering the unprotected system using

[2]The implementation of the *mobio-all* 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]The implementation of each face recognition model is available at https://gitlab.idiap.ch/bob/bob.bio.face

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

[6]Source code: https://gitlab.idiap.ch/bob/bob.paper.eusipco2023_mlphash

[7]The corresponding plots for other models are also available in the software package of the paper.

TABLE III: Comparison of MLP-Hash-protected, BioHash-protected, IoM-GRP-protected, IoM-URP-protected, and unprotected (Baseline) SOTA Face Recognition models, in terms of TMR (%) in the *normal* and the *stolen* scenarios on the MOBIO and LFW datasets. The threshold in each system is selected individually at an FMR of 10^{-3}. The results are reported as (mean±std) for 10 different experimental trials.

Dataset	Model	Baseline	*normal* scenario				*stolen* scenario			
			MLP-Hash	BioHash	IoM-GRP	IoM-URP	MLP-Hash	BioHash	IoM-GRP	IoM-URP
MOBIO	ArcFace	100.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	99.59 ± 0.08	100.00 ± 0.00	99.95 ± 0.04	99.98 ± 0.03	98.88 ± 0.13
	FaceNet	97.87	99.05 ± 0.48	99.93 ± 0.04	99.99 ± 0.01	95.56 ± 0.50	76.40 ± 6.19	89.38 ± 2.12	94.41 ± 0.83	87.51 ± 1.03
	IncResNetV2	96.69	99.98 ± 0.04	99.99 ± 0.01	100.00 ± 0.00	99.96 ± 0.03	65.12 ± 5.07	76.46 ± 7.49	92.56 ± 2.73	91.38 ± 1.35
LFW	ArcFace	98.73	98.86 ± 0.13	98.84 ± 0.05	99.19 ± 0.06	88.78 ± 1.37	95.54 ± 0.54	98.56 ± 0.06	98.62 ± 0.06	84.79 ± 1.98
	FaceNet	93.17	90.90 ± 0.90	96.81 ± 1.12	99.38 ± 0.09	69.29 ± 2.99	59.42 ± 5.02	83.19 ± 5.32	85.78 ± 4.92	50.77 ± 6.96
	IncResNetV2	93.33	99.42 ± 0.44	99.95 ± 0.05	100.00 ± 0.00	98.06 ± 0.35	47.40 ± 14.22	64.13 ± 19.89	83.38 ± 6.27	81.38 ± 2.75

the inverted templates. Hence, a higher SAR shows that the templates are more invertible, while a lower (or zero) SAR indicates that the protected templates are harder to invert.

To evaluate such an attack, similar to [26], we used a numerical solver (implemented in the SciPy package[8]) to find an estimate of the original template, which is mapped to the same output through the template protection module. The solver starts from an initial guess, and through an iterative process, tries to find an answer which gives the same output (as the given protected template) when passed as the input to the MLP-Hash with the same key. We also assumed that the attacker knows the distribution of unprotected templates, and uses this distribution to extract 10 samples as initial guesses in separate attempts. In each attempt, in the case of convergence of the solver, the inverted template is used to enter an unprotected system with a match threshold at a False Match Rate (FMR) of 10^{-3} (using the same feature extraction module).

Table II compares the irreversibility of MLP-Hash, BioHash, IoM-GRP, and IoM-URP protected templates of the ArcFace embeddings on the MOBIO database in terms of the SAR. As this table shows, the irreversibility of MLP-Hash is comparable to that of the BioHash and IoM-GRP methods. However, IoM-URP protected templates are more difficult to invert using our adopted inversion technique.

D. Recognition Accuracy Evaluation

To evaluate the recognition accuracy of MLP-Hash, we considered two scenarios: the *normal* scenario and the *stolen token* scenario. In the *normal* scenario, which is the expected scenario for most cases, each user's key is assumed to be secret. However, in the *stolen token* scenario (or briefly *stolen* scenario), we assume that the impostor has access to the user's secret key and uses this key with the impostor's own unprotected template. To implement the *stolen* scenario, in the verification stage we used the same key as the genuine's key for other users in the database to generate their MLP-Hash templates.

Table III compares the MLP-Hash-protected, BioHash-protected, IoM-GRP-protected, IoM-URP-protected, and unprotected (baseline) templates of the SOTA face recognition models, in terms of True Match Rate (TMR) in the *normal* and the *stolen* scenarios on the MOBIO and LFW datasets. The threshold in each system is selected individually at an FMR of 10^{-3}. As this table shows, in the normal scenario, all the protection schemes achieve comparable performance on the MOBIO dataset. However, on the LFW dataset, IoM-URP clearly has the worst performance. In the stolen scenario, IoM-GRP appears to perform the best across all three face recognition models and both evaluation datasets.

[8]https://scipy.org/

TABLE IV: Complexity comparison of template protection methods in terms of average execution time (milliseconds). The results are reported as (mean±std) for 1000 different experimental trials.

MLP-Hash	BioHash	IoM-GRP	IoM-URP
61.9 ± 0.5	12.5 ± 0.5	77.6 ± 0.2	36.2 ± 0.9

E. Discussion

Table I, Table II, and Table III compare the unlinkability, irreversibility and recognition accuracy, respectively, of our proposed template protection method with the BioHash, IoM-GRP, and IoM-URP algorithms. Table IV also compares the complexity of the aforementioned methods in generating protected templates from the ArcFace model in terms of average execution time (milliseconds) on a system equipped with an Intel(R) Core(TM) i7-7700K CPU @ 4.20GHz. Based on these results, IoM-URP is the most irreversible algorithm, however it clearly has the worst performance in the normal scenario (which is the expected scenario in practice). IoM-GRP has slightly better irreversibility than MLP-Hash, and its recognition accuracy is the best in most cases. However, it has the longest execution time amongst the studied protection methods. BioHashing has comparable recognition accuracy with MLP-Hash, and has slightly worse irreversibility. However, BioHashing has the shortest execution time. All in all, our experiments show that while all these template protection schemes have comparable unlinkability, there is a trade-off between irreversibility, recognition accuracy, and complexity.

IV. Conclusion

In this paper, we proposed a new cancelable biometric template protection scheme, dubbed MLP-Hash, which

uses a user-specific randomly-weighted multi-layer perceptron (MLP) with non-linear activation functions, followed by binarization of the output. We evaluated the unlinkability, irreversibility and recognition accuracy of MLP-Hash as per the ISO/IEC 30136 standard requirements, using SOTA face recognition models. Our protection method was found to satisfy these criteria to a high degree. In addition, we compared MLP-Hash with the BioHashing and IoM Hashing (IoM-GRP and IoM-URP) protection algorithms on the same SOTA face recognition systems, in terms of the recognition accuracy, unlinkability, and irreversibility criteria. Our experiments indicate that while all these template protection schemes are almost unlinkable, there is a trade-off between irreversibility, recognition accuracy, and complexity.

REFERENCES

[1] Guangcan Mai, Kai Cao, Pong C Yuen, and Anil K Jain, "On the reconstruction of face images from deep face templates," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 5, pp. 1188–1202, 2018.

[2] Hatef Otroshi Shahreza, Vedrana Krivokuća Hahn, and Sébastien 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, pp. 1211–1215.

[3] General Data Protection Regulation, "Regulation eu 2016/679 of the european parliament and of the council of 27 april 2016," *Official Journal of the European Union*, 2016.

[4] Karthik Nandakumar and Anil K Jain, "Biometric template protection: Bridging the performance gap between theory and practice," *IEEE Signal Processing Magazine*, vol. 32, no. 5, pp. 88–100, 2015.

[5] Arpita Sarkar and Binod K Singh, "A review on performance, security and various biometric template protection schemes for biometric authentication systems," *Multimedia Tools and Applications*, pp. 1–56, 2020.

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

[7] Andrew Teoh Beng Jin, David Ngo Chek Ling, and Alwyn Goh, "Biohashing: two factor authentication featuring fingerprint data and tokenised random number," *Pattern Recognition*, vol. 37, no. 11, pp. 2245–2255, 2004.

[8] Zhe Jin, Jung Yeon Hwang, Yen-Lung Lai, Soohyung Kim, and Andrew Beng Jin Teoh, "Ranking-based locality sensitive hashing-enabled cancelable biometrics: Index-of-max hashing," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 2, pp. 393–407, 2017.

[9] Mulagala Sandhya and Munaga VNK Prasad, "Biometric template protection: A systematic literature review of approaches and modalities," in *Biometric Security and Privacy*, pp. 323–370. Springer, 2017.

[10] Ari Juels and Martin Wattenberg, "A fuzzy commitment scheme," in *Proceedings of the 6th ACM Conference on Computer and Communications Security*, 1999, pp. 28–36.

[11] Ari Juels and Madhu Sudan, "A fuzzy vault scheme," *Designs, Codes and Cryptography*, vol. 38, no. 2, pp. 237–257, 2006.

[12] Umut Uludag, Sharath Pankanti, Salil Prabhakar, and Anil K Jain, "Biometric cryptosystems: issues and challenges," *Proceedings of the IEEE*, vol. 92, no. 6, pp. 948–960, 2004.

[13] Jiankang Deng, Jia Guo, Xue Niannan, and Stefanos Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[14] Florian Schroff, Dmitry Kalenichenko, and James Philbin, "Facenet: A unified embedding for face recognition and clustering," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 815–823.

[15] Tiago de Freitas Pereira, André Anjos, and Sébastien Marcel, "Heterogeneous face recognition using domain specific units," *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 7, pp. 1803–1816, 2018.

[16] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller, "Labeled faces in the wild: A database for studying face recognition in unconstrained environments," Tech. Rep. 07-49, University of Massachusetts, Amherst, October 2007.

[17] Chris McCool, Roy Wallace, Mitchell McLaren, Laurent El Shafey, and Sébastien Marcel, "Session variability modelling for face authentication," *IET Biometrics*, vol. 2, no. 3, pp. 117–129, Sept. 2013.

[18] Guangcan Mai, Kai Cao, Xiangyuan Lan, and Pong C Yuen, "Secureface: Face template protection," *IEEE Transactions on Information Forensics and security*, vol. 16, pp. 262–277, 2020.

[19] Young Kyun Jang and Nam Ik Cho, "Deep face image retrieval for cancelable biometric authentication," in *2019 16th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*. IEEE, 2019, pp. 1–8.

[20] Hakyoung Lee, Cheng Yaw Low, and Andrew Beng Jin Teoh, "Softmax-out transformation-permutation network for facial template protection," in *2020 25th International Conference on Pattern Recognition (ICPR)*. IEEE, 2021, pp. 7558–7565.

[21] A. Anjos, L. El 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.

[22] 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.

[23] Hatef Otroshi Shahreza and Sébastien Marcel, "Towards protecting and enhancing vascular biometric recognition methods via biohashing and deep neural networks," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 3, no. 3, pp. 394–404, 2021.

[24] Hatef Otroshi Shahreza, Vedrana Krivokuća Hahn, and Sébastien Marcel, "On the recognition performance of biohashing on state-of-the-art face recognition models," in *2021 IEEE International Workshop on Information Forensics and Security (WIFS)*. IEEE, 2021, pp. 1–6.

[25] Marta Gomez-Barrero, Javier Galbally, Christian Rathgeb, and Christoph Busch, "General framework to evaluate unlinkability in biometric template protection systems," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 6, pp. 1406–1420, 2017.

[26] Vedrana Krivokuća Hahn and Sébastien Marcel, "Towards protecting face embeddings in mobile face verification scenarios," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, pp. 1–1, 2022.