Efficiency Analysis of Post-quantum-secure Face Template Protection Schemes based on Homomorphic Encryption

Jascha Kolberg[1], Pawel Drozdowski[1], Marta Gomez-Barrero[2], Christian Rathgeb[1], Christoph Busch[1]

Abstract: Since biometric characteristics are not revocable and biometric data is sensitive, privacy-preserving methods are essential to operate a biometric recognition system. More precisely, the biometric information protection standard ISO/IEC IS 24745 requires that biometric templates are stored and compared in a secure domain. Using homomorphic encryption (HE), we can ensure permanent protection since mathematical operations on the ciphertexts directly correspond to those on the plaintexts. Thus, HE allows to compute the distance between two protected templates in the encrypted domain without a degradation of biometric performance with respect to the corresponding system. In this paper, we benchmark three post-quantum-secure HE schemes, and thereby show that a face verification in the encrypted domain requires only 50 ms transaction time and a template size of 5.5 KB.

Keywords: Face Recognition, Biometric Template Protection, Post-quantum Cryptography, Homomorphic Encryption.

1 Introduction

Nowadays, biometric authentication is widely used in applications ranging from convenient smartphone unlocking to high-security border control. On the other hand, we can also observe an increase in cybercrime and databases leakages. Due to the fact that biometric characteristics are unique and cannot be changed unlike e.g. passwords, unprotected databases can be exploited to reveal enrolment data and track individuals. Hence, biometric data, amongst others, are classified as sensitive data by the European Union in the General Data Protection Regulation 2016/679 [EU16]. Furthermore, research has proven that biometric samples can be reconstructed from unprotected templates, for instance face [Ma18], iris [Ga13], or fingerprint [Ca07]. The ISO/IEC IS 24745 standard [IS11] defines three requirements for biometric template protection (BTP): *i) unlinkability*, two protected templates cannot be linked to the same subject, *ii) renewability*, new templates can be created without the need to re-enrol and old templates can be revoked, and *iii) irreversibility*, it is impossible to retrieve original samples given only protected templates. Furthermore, the biometric performance should be preserved in the protected scheme. Therefore, BTP mechanisms are able to handle these privacy issues since templates are stored and compared in a secure domain.

[1] da/sec - Biometrics and Internet Security, University of Applied Sciences Darmstadt, Germany, {jascha.kolberg,pawel.drozdowski,christian.rathgeb,christoph.busch}@h-da.de

[2] Hochschule Ansbach, Germany, marta.gomez-barrero@hs-ansbach.de

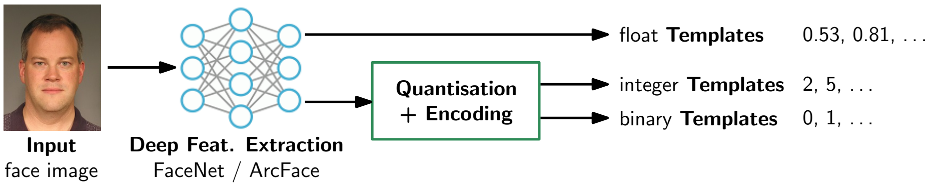
Input
face image
Deep Feat. Extraction
FaceNet / ArcFace
Quantisation
+ Encoding
float Templates 0.53, 0.81, ...
integer Templates 2, 5, ...
binary Templates 0, 1, ...

Fig. 1: Pre-processing pipeline in order to extract templates from facial input images.

By fulfilling the aforementioned requirements, the data subject's privacy is protected during the comparison as well as for leaked templates. In this context, different BTP schemes have been developed [BDL15, RU11]. One trend utilises homomorphic encryption (HE) in order to compute the biometric comparison score directly on the ciphertexts. For instance, Sadeghi *et al.* [SSW09] combine HE with garbled circuits to achieve protection in a face identification scenario. Using two HE schemes for face identification, Drozdowski *et al.* [Dr19] additionally discuss technical considerations and challenges. In order to speed up the execution for face verifications, Boddeti [Bo18] explores fully HE in combination with batching, which allows to reduce the number of homomorphic multiplications for the distance calculation. Likewise, Yasuda *et al.* [Ya15] present a specific packing method for their HE scheme to gain efficiency. Therewith, they apply template protection for different biometric modalities with feature vectors of 2,048 bit. Following the multi-modal idea, Gomez-Barrero *et al.* [Go17] propose a general framework for verification of fused modalities in the homomorphically encrypted domain.

However, a lot of publications applying HE for BTP assign the secret key to the client, thus discarding the advantages of using biometrics in general and generating a two-factor authentication system with biometrics and a secret-based knowledge. In contrast, our contribution keeps the secret key on the server side. We build upon the iris BTP scheme in [Ko19] to protect binarised face templates and benchmark recognition performance, transaction time, template size, and cryptographic security with two state-of-the-art HE BTP systems [Dr19]. The focus is on post-quantum-secure cryptography [BL17] to achieve long-term security for biometric data.

The rest of this paper is structured as follows: Section 2 describes the proposed system including the necessary pre-processing. The experimental setup and results, including the benchmark, are presented in Section 3. Finally, Section 4 concludes the paper.

2 Proposed System

2.1 Baseline system

Before we can encrypt face templates, we need to extract the features from the input images. The pre-processing pipeline for this purpose is shown in Fig. 1. Given the facial input images, two deep feature extraction algorithms, ArcFace [De19] and FaceNet [SKP15], were used to create templates of 512 floating-point values. Additionally, we applied quantisation and encoding to transform the float templates into integer and binary templates in order to be able to use additional HE schemes. Following the analysis in [Dr18], the feature space is divided into four segments of equal size. For the integer encoding, the float values are simply mapped to the corresponding number of their sequence area. In order to have the lowest distance for adjacent areas in the binary representation, the linearly separable subcode (LSSC) [LT12] transforms each integer value into three binary digits. Float and integer templates can be compared by computing the squared Euclidean distance, while the Hamming distance is used on binary templates.

2.2 Homomorphic encryption

Homomorphic encryption schemes [Ac18] implement asymmetric cryptography with the property that specific mathematical operations on the ciphertext directly affect to the plaintext. Those additive or multiplicative homomorphic properties can generally be defined as:

$$\mathrm{Enc}\,(A+B) = \mathrm{Enc}\,(A) \diamond \mathrm{Enc}\,(B) \tag{1}$$

$$\mathrm{Enc}\,(A \cdot B) = \mathrm{Enc}\,(A) \circ \mathrm{Enc}\,(B) \tag{2}$$

Hence, we have an operation $\diamond$ that results in the sum of two plaintexts when it is applied to both corresponding ciphertexts. Another operation $\circ$ is used to achieve a multiplication. The specific operations depend on the selected HE scheme. Moreover, not all HE schemes support all operations. Therefore, depending on the used biometric templates and the required distance computations, different HE schemes [Ac18] should be utilised.

With the focus on post-quantum-security [BL17], the following three crypto schemes are selected. In order to compute the squared Euclidean distance in the encrypted domain, two HE schemes are utilised. The encryption scheme by Cheon-Kim-Kim-Song (CKKS) [Ch17] supports homomorphic operations on floating point templates and the Brakerski/Fan-Vercauteren (BFV) [FV12] scheme is applied on the integer templates. On the other hand, the computation of the Hamming distance in the encrypted domain can efficiently be done with the N-th degree truncated polynomial ring (NTRU) [HPS98] scheme, if the parameters are selected in a way that the decryption automatically performs a modulo-2 operation. Then one addition directly results in the *XOR* of probe and reference. The security of all three schemes is based on the ring-learning-with-errors problem, which, using a quantum algorithm, can be reduced to the shortest vector problem over ideal lattices [LPR10]. Thus, granting long-term security for biometric templates.

2.3 Protected system

Based on the aforementioned considerations, we can build our HE protected system. Since we are not interested in a two-factor authentication, where the secret key is assigned to the client, the secret key needs to be stored at server-side. However, placing the decryption key next to the encrypted templates in the database server (DB) threatens the whole purpose. Hence, we need an additional authentication server (AS) in our infrastructure, which works as a trusted third party. The structure of this system and a verification transaction are illustrated in Fig. 2.

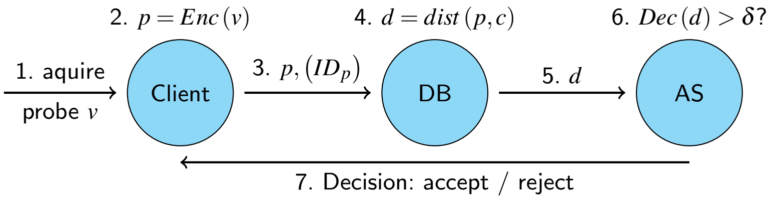

Fig. 2: Verification transaction of the BTP system using homomorphic encryption.

1. The client captures the biometric characteristics and pre-processes the data, resulting in a probe feature vector v.
2. The client encrypts v with the public key to get the protected probe p.
3. The encrypted probe p is sent to the database server (DB). In a verification scenario, the client additionally transfers an ID claim.
4. DB computes the distance d between probe p and reference(s) c_i in the encrypted domain.
5. This encrypted distance d is forwarded to the authentication server (AS).
6. AS decrypts d using the secret key and compares the result with a decision threshold. Alternatively, AS could also sort all computed distances in identification mode.
7. The final accept/reject decision is revealed to the client.

This architecture assumes the honest-but-curious model, where the parties stick to the protocol, but may try to learn as much information as possible. This implies that DB and AS do not collude in order to decrypt the database or incoming probes. The client encrypts its probe before sending it to the DB, which only operates on encrypted templates to compute the encrypted distance. The AS, which possesses the secret decryption key, receives only protected distance values and thus does not learn sensitive information from neither the probe nor the reference. The transmission channel between parties can additionally be protected by TLS. For higher privacy, the decision (in 7.) could be returned to the DB, which forwards it to the client in order to conceal the identity of the client device from AS.

3 Experimental Evaluation

3.1 Experimental Setup

The experiments were run on a frontal image subset of the FERET database [Ph00] comprising 6,963 samples of 563 subjects. Both feature extraction methods use their freely available pre-trained model, which allows for reproducibility of our research. Furthermore, our BTP systems are implemented based on the open-source crypto implementations in the Microsoft SEAL HE library[3] (CKKS and BFV) and the NTRU iris template protection system [Ko19].

[3] `https://github.com/Microsoft/SEAL`

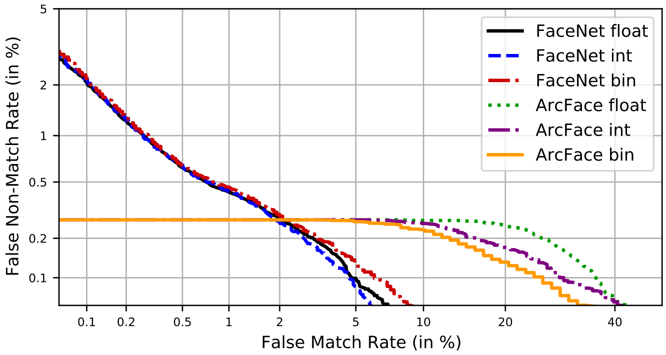

Fig. 3: DET plot showing the false match rates and false non-match rates for different template types in the verification scenario. The performance is identical for unprotected and protected systems.

Tab. 1: Rank-1 identification rates in % for different feature types.

Rank-1	float	integer	binary
ArcFace (%)	99.03	98.98	99.03
FaceNet (%)	98.50	98.42	98.36

In order to compare the biometric performance of different template representations, verification scores as well as rank-1 mated identification scores are computed. For the genuine verification, all mated samples are compared and for the impostor scores, only the first sample of each subject is compared with the first sample of all other subjects. Furthermore, the rank-1 identification scores include all mated comparisons.

The timing of relevant functions was conducted within a virtualised (single-core) Linux on a commodity notebook running an Intel Core i7 2.7 GHz CPU and 16 GB DDR4 RAM. While the SEAL library is written in C++, the Python3 code of the NTRU implementation is executed with PyPy3[4] for an additional speed-up.

3.2 Results

Biometric performance evaluation. Fig. 3 shows that biometric verification performance is preserved across the different template types. Especially for a false match rate below 2%, the DET curves for all types are almost identical. However, looking at the rank-1 identification rates of all mated comparisons in Tab. 1 reveals minor variations for the different template types due to the quantisations. In general we can conclude, that the biometric performance remains stable across different face template representations.

Execution time and file size. Generally, we could observe that execution times and file sizes slightly increase for higher security levels. However, the relative speed-up between different template representations stayed the same and hence only results for a security level of 128 bits are presented.

[4] https://pypy.org

Tab. 2: Median execution time and standard deviation of relevant functions. The comparison includes distance computation, decryption, and deriving the final decision.

128 bits Security	CKKS (float)	BFV (int)	NTRU (bin)
Key generation (ms)	779 (±4)	255 (±5)	362 (±84)
Encryption (ms)	6 (±2)	76 (±1)	27 (±5)
Comparison (ms)	3391 (±10)	618 (±26)	23 (±3)

Tab. 3: File size of keys and a single template for the evaluated encryption schemes.

128 bits Security	CKKS (float)	BFV (int)	NTRU (bin)
Keys	99 MB	12 MB	6 KB
Template	516 KB	132 KB	5.5 KB

The execution times for relevant functions are depicted in Tab. 2. The key generation of all encryption schemes is done in less than one second and is negligible since it is a one time effort at enrolment. The encryption times refer to a single template and need to be multiplied with the number of references in the database during the system setup. An encryption takes about 6 ms for float features, 76 ms for integer features, and 27 ms for binary features. Analogously to the encryption, we need to multiply the comparison times with the number of references for each performed identification. A comparison in BFV (618 ms) is five times faster than CKKS (3,391 ms) and 25 times slower than NTRU (23 ms). Additionally, for both identification and verification, the probe needs to be encrypted before the comparison. Thus, this single encryption affects the verification much more than the identification scenario. It occurs that the CKKS encryption is faster than the other encryptions, while it needs much longer on other operations. The NTRU encryption relies on new random polynomials for each block and their generation takes apparently more time than its corresponding function within CKKS. Comparing the file sizes, as shown in Tab. 3, reveals intuitive results. The keys of CKKS are with 99 MB the biggest, while BFV requires 12 MB and NTRU 6 KB. The same order holds for the template sizes; 516 KB for CKKS, 132 KB for BFV, and 5.5 KB for NTRU.

Simulating a company database with 1,000 employees would require a database storage of around 500 MB for CKKS, 130 MB for BFV, and 6 MB for NTRU. Using the same system in identification mode would take around one hour in the CKKS scheme, nearly 12 minutes in the BFV scheme, and 50 seconds with NTRU.

Security analysis. Finally, all three HE schemes are based on ideal lattices and hence found to be post-quantum-secure [BL17], which grants us *irreversibility*. The encryption functions utilise a random factor with the effect that encrypting the identical plaintext twice, results in two *unlinkable* ciphertexts. *Renewability* can be achieved by exchanging the key pair and re-encrypting the database. Since clients only operate with the public key, no re-enrolment is required.

Summary. These results show that real time verifications in the encrypted domain are possible for integer and binary face templates. However, when it comes to identification, only the comparison of binary templates is fast enough to support a reasonable transaction time. As demonstrated, the biometric performance remains largely unaffected for all presented template representations.

4 Conclusions

This work showed that the most important requirement for efficient template protection, transforming float templates into integer or binary templates, has negligible impact on the biometric recognition accuracy. Furthermore, the three ISO/IEC IS 24745 requirements *irreversibility*, *unlinkability*, and *renewability* are fulfilled by the evaluated HE schemes, CKKS, BFV, and NTRU. Additional long-term template security is granted by their post-quantum-secure design. Since we worked with a public database and only used open-source software, all results from this paper are reproducible. Most importantly, using binary face templates, a verification in the encrypted domain is done within 50 ms on an ordinary notebook, which also allows to apply NTRU HE in limited identification scenarios. Those results demonstrate the practicability of biometric template protection for face verification even on off-the-shelf hardware. Future work will evaluate efficient biometric identification in the homomorphic domain including computational workload reduction methods [DRB19].

Acknowledgements

This research work has been funded by the German Federal Ministry of Education and Research and the Hessen State Ministry for Higher Education, Research and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE.

References

[Ac18] Acar, A.; Aksu, H.; Uluagac, A. S.; Conti, M.: A Survey on Homomorphic Encryption Schemes: Theory and Implementation. ACM Computing Surveys, 51(4):1–35, 2018.

[BDL15] Barni, M.; Droandi, G.; Lazzeretti, R.: Privacy Protection in Biometric-Based Recognition Systems: A Marriage Between Cryptography and Signal Processing. IEEE Signal Processing Magazine, 32(5):66–76, 2015.

[BL17] Bernstein, D. J.; Lange, T.: Post-quantum cryptography. Nature, 549(7671):188–194, 2017.

[Bo18] Boddeti, V. N.: Secure Face Matching Using Fully Homomorphic Encryption. In: Proc. of Int. Conf. on Biometrics Theory, Applications and Systems (BTAS). IEEE, pp. 1–10, 2018.

[Ca07] Cappelli, R.; Maio, D.; Lumini, A.; Maltoni, D.: Fingerprint Image Reconstruction From Standard Templates. Trans. on Pattern Analysis and Machine Intelligence, 29(9), 2007.

[Ch17] Cheon, J. H.; Kim, A.; Kim, M.; Song, Y.: Homomorphic Encryption for Arithmetic of Approximate Numbers. In: Proc. Asiacrypt. Springer, pp. 409–437, 2017.

[De19] Deng, J.; Guo, J.; Xue, N.; Zafeiriou, S.: ArcFace: Additive Angular Margin Loss for Deep Face Recognition. In: Proc. CVPR. pp. 4690–4699, 2019.

[Dr18] Drozdowski, P.; Struck, F.; Rathgeb, C.; Busch, C.: Benchmarking Binarisation Schemes for Deep Face Templates. In: Proc. ICIP. IEEE, pp. 1–5, 2018.

[Dr19] Drozdowski, P.; Buchmann, N.; Rathgeb, C.; Margraf, M.; Busch, C.: On the Application of Homomorphic Encryption to Face Identification. In: Proc. BIOSIG. pp. 1–8, 2019.

[DRB19] Drozdowski, P.; Rathgeb, C.; Busch, C.: Computational Workload in Biometric Identification Systems: An Overview. IET Biometrics, 8(6):351–368, 2019.

[EU16] EU Parliament: . EU Regulation 2016/679 (General Data Protection Regulation), 2016.

[FV12] Fan, J.; Vercauteren, F.: Somewhat Practical Fully Homomorphic Encryption. IACR Cryptology ePrint Archive, 2012:144, 2012.

[Ga13] Galbally, J.; Ross, A.; Gomez-Barrero, M.; Fierrez, J. et al.: Iris Image Reconstruction From Binary Templates: An Efficient Probabilistic Approach Based on Genetic Algorithms. Computer Vision and Image Understanding, 117(10):1512–1525, 2013.

[Go17] Gomez-Barrero, M.; Maiorana, E.; Galbally, J.; Campisi, P.; Fierrez, J.: Multi-Biometric Template Protection Based on Homomorphic Encryption. Pattern Recognition, 67:149–163, 2017.

[HPS98] Hoffstein, J.; Pipher, J.; Silverman, J. H.: NTRU: A Ring-Based Public Key Cryptosystem. In: Int. Algorithmic Number Theory Symposium. Springer, pp. 267–288, 1998.

[IS11] ISO/IEC JTC1 SC27 Security Techniques: . ISO/IEC 24745:2011. Information Technology - Security Techniques - Biometric Information Protection. ISO/IEC, 2011.

[Ko19] Kolberg, J.; Bauspieß, P.; Gomez-Barrero, M.; Rathgeb, C.; Dürmuth, M.; Busch, C.: Template Protection based on Homomorphic Encryption: Computationally Efficient Application to Iris-Biometric Verification and Identification. In: Proc. WIFS. 2019.

[LPR10] Lyubashevsky, V.; Peikert, C.; Regev, O.: On Ideal Lattices and Learning with Errors Over Rings. In: Annual Int. Conf. on the Theory and Applications of Cryptographic Techniques. Springer, pp. 1–23, 2010.

[LT12] Lim, M. H.; Teoh, A. B. J.: A Novel Encoding Scheme for Effective Biometric Discretization: Linearly Separable Subcode. IEEE Trans. on Pattern Analysis and Machine Intelligence, 35(2):300–313, 2012.

[Ma18] Mai, G.; Cao, K.; Yuen, P. C.; Jain, A. K.: On the Reconstruction of Face Images from Deep Face Templates. IEEE Trans. on Pattern Analysis and Machine Intelligence, 2018.

[Ph00] Phillips, P. J.; Moon, H.; Rizvi, S. A.; Rauss, P. J.: The FERET Evaluation Methodology for Face-Recognition Algorithms. IEEE Trans. on Pattern Analysis and Machine Intelligence, 22(10):1090–1104, 2000.

[RU11] Rathgeb, C.; Uhl, A.: A Survey on Biometric Cryptosystems and Cancelable Biometrics. EURASIP Journal on Information Security, 2011(1):3, 2011.

[SKP15] Schroff, F.; Kalenichenko, D.; Philbin, J.: FaceNet: A Unified Embedding for Face Recognition and Clustering. In: Proc. CVPR. pp. 815–823, 2015.

[SSW09] Sadeghi, A. R.; Schneider, T.; Wehrenberg, I.: Efficient Privacy-preserving Face Recognition. In: Proc. Inscrypt. Springer, pp. 229–244, 2009.

[Ya15] Yasuda, M.; Shimoyama, T.; Kogure, J.; Yokoyama, K.; Koshiba, T.: New Packing Method in Somewhat Homomorphic Encryption and Its Applications. Security and Communication Networks, 8(13):2194–2213, 2015.