

On the Application of Homomorphic Encryption to Face Identification

P. Drozdowski^{1,3} N. Buchmann² C. Rathgeb³ M. Margraf² C. Busch³

Abstract: The data security and privacy of enrolled subjects is a critical requirement expected from biometric systems. This paper addresses said topic in facial biometric identification. In order to fulfil the properties of unlinkability, irreversibility, and renewability of the templates required for biometric template protection schemes, homomorphic encryption is utilised. In addition to achieving the aforementioned objectives, the use of homomorphic encryption ensures that the biometric performance remains completely unaffected by the template protection scheme.

The main contributions of this paper are: It proposes an architecture of a system capable of performing biometric identification in the encrypted domain, as well as provides and evaluates an implementation using two existing homomorphic encryption schemes. Furthermore, it discusses the pertinent technical considerations and challenges in this context.

Keywords: Biometric Identification, Template Protection, Face, Homomorphic Encryption

1 Introduction

Data exposure is a potential risk in biometric system deployments, which typically store their data secured using traditional encryption algorithms [NJ15]. If this protection were to be compromised, serious problems would arise, including but not limited to, identity theft, cross-matching without consent, and severely limited renewability. Increasingly, the public and non-governmental organisations pay attention to those (and other) issues associated with centralised storage of sensitive personal and biometric data. In some areas, this contributes to the political process of widening legislation against privacy violations (*e.g.* GDPR in Europe [Eu16]), which entail significant responsibilities for the data controllers.

Recently, biometric *template protection* (see *e.g.* [RU11] for a survey) has been an active research field attempting to address said security and privacy challenges. The ISO/IEC IS 24745 [IS11] mandates several properties, which must be guaranteed by such schemes:

Unlinkability referring to making it infeasible to determine if two or more protected templates were derived from the same instance. By fulfilling this property, cross-matching across different databases is prevented.

Irreversibility referring to making it infeasible to reconstruct the original biometric data given a protected template and its corresponding secret. With this property fulfilled, the privacy of the users' data is increased, and additionally the security of the system is increased against presentation and replay attacks.

Renewability referring to making it possible to revoke old protected templates and creating new ones from the same biometric instance and/or sample. With this property fulfilled, it is possible to revoke and reissue the templates in case of the database being compromised, thereby preventing misuse.

¹ Norwegian Biometrics Laboratory, NTNU, Gjøvik, Norway

² Freie Universität Berlin, Germany
{nicolas.buchmann,marian.margraf}@fu-berlin.de

³ da/sec – Biometrics and Internet Security Research Group, Hochschule Darmstadt, Germany
{pawel.drozdowski,christian.rathgeb,christoph.busch}@h-da.de

Performance preservation referring to the requirement of the biometric performance not being significantly impaired by the protection scheme.

Three main biometric template protection approach classes can be distinguished: (1) biometric cryptosystems, which use the biometric data to bind or extract a key [Ca13], (2) cancelable biometrics, which utilise irreversible transformations to the biometric samples or templates [PRC15], and (3) (homomorphic) encryption of biometric data [Ag13].

Homomorphic encryption (henceforth referred to as “HE”) makes it possible to compute operations in the encrypted domain, which render the same result as those in the plaintext domain. Thus, provided that it is possible to implement a given biometric comparator to feasibly operate within the homomorphic domain, such a template protection scheme would operate without any loss of biometric performance, whereas some impairment is often inevitable in biometric cryptosystems and cancelable biometrics. In general, an encryption algorithm E has the homomorphic property for an operation \odot if it holds $E(m_1) \odot E(m_2) = E(m_1 \odot m_2), \forall m_1, m_2 \in M$ where M is the set of all possible messages. HE schemes are classified depending on the number and type of supported \odot operations (see *e.g.* [Ac18] for a detailed survey). The following three HE scheme types exist today:

Partially Homomorphic Encryption (PHE) schemes are defined as allowing only a single operation type an unlimited number of times. PHE schemes have been around for over 30 years and are the oldest HE schemes like the classical RSA scheme [RSA78] and the El-Gamal scheme [EG85] supporting only either addition or multiplication.

Somewhat Homomorphic Encryption (SWHE) schemes allow multiple operation types, but only a limited number of times. SWHE examples from literature are Yao’s garbled circuit scheme [Ya82], which supports arbitrary operations a limited number of times and the Boneh-Goh-Nissim (BGN) scheme [BGN05], which supports unlimited number of additions and one multiplication.

Fully Homomorphic Encryption (FHE) schemes support an unlimited number of operations. The first feasible FHE scheme was proposed by Gentry [Ge09] in 2009 and many newer FHE schemes are based on Gentry’s general FHE framework.

Brakerski and Vaikuntanathan [BV11] utilise Gentry’s framework to make their SWHE scheme a FHE scheme and introduce batching as an optimization [BGV14]. The Brakerski’s scheme was later optimised by Fan and Vercauteren [FV12].

While several authors investigated using FHE for biometric *verification* (see *e.g.* [Go17, ISR18, Bo19]) with promising results, the biometric *identification* scenario has not yet been addressed or merely considered a trivial extension of the proposed schemes. However, there exist several challenges and issues which must be dealt with for such schemes to be viable in the biometric identification scenario, especially if computational workload reduction (*i.e.* decreasing the computational complexity of the retrieval, see *e.g.* [KPB18] for a survey) is to be employed. Accordingly, the contribution of this paper is twofold: (1) an example architecture of a system capable of performing biometric identification with homomorphically protected templates is described, implemented, and evaluated with a facial recognition system, and (2) the practical considerations and challenges relevant to the biometric identification scenario are discussed in the context of HE and potential solutions, along with the future research avenues being explored.

This paper is organised as follows: in section 2 the proposed system is described. Section 3 outlines the experimental setup and the results of the evaluations. The results and other relevant matters are discussed in section 4, while concluding remarks and a summary are given in section 5.

2 Proposed System

Figure 1 shows an overview of the proposed system. There are 3 entities in the system: A client, where the biometric features are extracted (not depicted) and encrypted; a database, where the encrypted references of the enrolled subjects are stored and the distances between them and the probe computed in the encrypted domain and applies a decision threshold; and a trusted third party (TTP), which decrypts the thresholded scores and communicates a decision to the client. This description is an abstract one, the concrete HE schemes used in the implementation and evaluation are listed in section 3.

The proposed system ensures that the unencrypted, privacy-sensitive biometric features are only available to the client. All network transfers of the biometric probe, as well as the computations on the data happen in the encrypted domain, thereby preventing eavesdropping attacks on the network and malicious or rogue database system attacks. Since all database entries are stored encrypted, a database breach or insider threats yield no valuable attack vectors. To further strengthen the security of the proposed system, following measures are possible: 1) to curtail active attacks (*e.g.* Man-in-the-Middle or skimming) against the transferred feature vectors and the identification transaction decision, TLS can be deployed between the parties; 2) to prevent the attacker from conducting an analysis of a single entry or unsanctioned database audits, the order of the database entries could be randomly shuffled during or after each identification transaction. Both measures would have a negligible impact on the access time and no impact on the biometric performance.

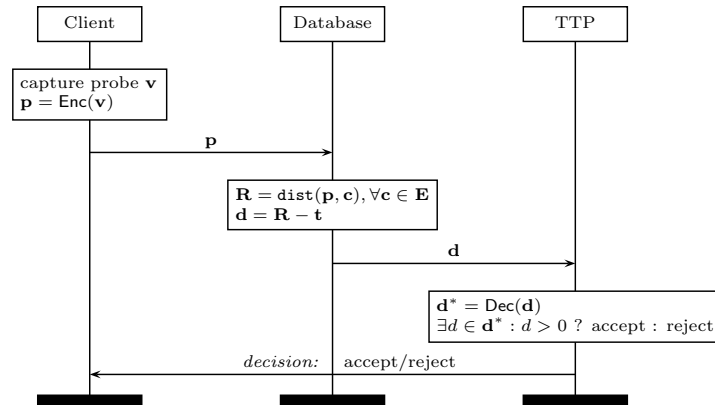


Fig. 1: System overview sketch

1. Let \mathbf{v} denote a biometric feature vector, with a constant number (n) of feature elements: $\mathbf{v} = (v_1, v_2, \dots, v_n)$ and \mathbf{c} denote the element-wise encrypted version of said reference vector using a HE scheme with batching. Batching is a mechanism which allows to perform operations on vectors rather than individual numbers in the encrypted domain, thereby allowing vast (by several orders of magnitude) speed-ups

on vectorisable operations. By design, the HE schemes used in this paper support a certain number of slots (s) for the elements in the encrypted vector. This number depends on the parameters of the encryption algorithm, including the encryption security. In this case (see subsection 3.1), $n < s$ for any reasonable (from the security perspective) combination of parameters, which means that \mathbf{c} must be padded with zeroes beyond the n -th index, *i.e.* $\mathbf{c} = (c_1, c_2, \dots, c_n, 0 \dots 0)$.

2. Let \mathbf{E} denote an enrolment database consisting of N such encrypted vectors (sub-

jects), *i.e.* $\mathbf{E} = \begin{bmatrix} \mathbf{c}_1 \\ \mathbf{c}_2 \\ \vdots \\ \mathbf{c}_N \end{bmatrix}$. Let \mathbf{p} denote an encrypted feature vector of a probe and \odot

an arbitrary function for distance computation between two biometric feature vectors. In this work, the squared Euclidean distance will be used, *i.e.* $\text{dist}(\mathbf{c}, \mathbf{p}) = \sum_{i=0}^s (\mathbf{c}_i - \mathbf{p}_i)^2$. While performing the subtraction and exponentiation in the encrypted domain is done trivially by directly using the operators in the encrypted domain, the summing up of the individual elements is not as straightforward, since batching techniques do not allow access to individual encrypted vector elements. To compute the sum, the observation made in [GH11] is utilised – the vector is successively circularly shifted and added onto itself s times. Afterwards, the sum of elements is present in the first element of the vector. The other elements of the vector are now irrelevant and are cleaned by multiplying with 0. Mathematically, a vector $\mathbf{k} = \{1, 0, \dots, 0\}$ is defined and multiplied with the result.

3. The goal of a biometric identification is to first compute the comparison scores between the probe and all the items in the enrolment database, select the best one, and compare it against a predetermined threshold to make the final decision. Let \mathbf{r}_i denote the result of comparison between the probe and the i 'th entry in the database, *i.e.* $\mathbf{r}_i = \text{dist}(\mathbf{p}, \mathbf{c}_i) \cdot \mathbf{k}$. After processing the whole enrolment database, N such result

vectors as separate ciphertexts are created: $\mathbf{R} = \begin{bmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \\ \vdots \\ \mathbf{r}_N \end{bmatrix} = \begin{bmatrix} (\mathbf{r}_{1,1}, 0, \dots, 0) \\ (\mathbf{r}_{2,1}, 0, \dots, 0) \\ \vdots \\ (\mathbf{r}_{N,1}, 0, \dots, 0) \end{bmatrix}$ note, that

each of \mathbf{r}_i is encrypted in a separate ciphertext, rather than \mathbf{R} being encrypted in one ciphertext. In the next step, those individual ciphertexts are combined. At this point, it is also possible to randomly shuffle the order of the ciphertexts, thereby preventing the trusted third party from later learning which subject has been identified.

4. Through shifting the \mathbf{r}_i vectors, a structure conceptually akin to a diagonal matrix

is reached, *i.e.* $\mathbf{R} = \begin{bmatrix} (\mathbf{r}_{1,1}, 0, \dots, 0) \\ (0, \mathbf{r}_{2,1}, \dots, 0) \\ \vdots \\ (0, 0, \dots, \mathbf{r}_{N,1}) \end{bmatrix}$. Those vectors are combined by adding them

together, thereby producing a single encrypted vector holding the comparison scores of \mathbf{p} against \mathbf{E} , *i.e.* a mapping was introduced so that $\mathbf{R} \mapsto (\mathbf{r}_{1,1}, \mathbf{r}_{2,1}, \dots, \mathbf{r}_{N,1})$.

5. The next step is to compare the scores against a predefined threshold, *i.e.* to transform from the continuous spectrum of the comparison scores to the binary accept or reject decision. This is done by subtracting an encrypted vector storing at each element the threshold value (t), *i.e.* $\mathbf{t} = (t, t, \dots, t)$ of length s , thus producing a decision vector $\mathbf{d} = \mathbf{R} - \mathbf{t}$, which is subsequently transmitted to the trusted third party.
6. In the last step, \mathbf{d} is decrypted to determine the identification outcome (and communicate it to the client), *i.e.* decision = $\begin{cases} \text{accept} & \text{if } \exists d \in \mathbf{d} : d > 0 \\ \text{reject} & \text{if } \forall d \in \mathbf{d} : d \leq 0 \end{cases}$. If any of the elements in \mathbf{d} is above 0, it is necessarily because the corresponding score in \mathbf{R} has been greater than t , thus yielding an accept decision. Since due to batching access to individual elements in the encrypted vector is not possible, the whole vector must be decrypted by the trusted third party.

3 Experiments

3.1 Experimental Setup

The experimental evaluation was conducted on a frontal subset of the FERET database [Ph00] in an open-set identification scenario with 500 enrolled data subjects, using 10-fold cross-validation. The features from the images were extracted using FaceNet with a pre-trained model provided by its authors [SKP15]. FaceNet yields templates comprising of 512 floating-point feature elements; two such templates can then be compared using squared Euclidean distance. The open-source Microsoft SEAL HE library [Mi19] was utilised to implement the identification protocol in the encrypted domain (see section 2). The choice was based on the presence of a high-level API and suitable HE schemes, namely: Brakerski/Fan-Vercauteren (henceforth referred to as ‘‘BFV’’) [FV12] and Cheon-Kim-Kim-Song (henceforth referred to as ‘‘CKKS’’) [Ch17] for integer and float based computations, respectively (see [Ac18] for a detailed HE survey, including other available libraries). Thus, the templates produced by FaceNet can be encrypted directly using CKKS. To utilise BFV, a quantisation scheme is employed, whereby the continuously distributed values of the feature elements are mapped into discrete intervals (see *e.g.* [Dr18] for more details). Although some information is lost through quantisation, the biometric performance should remain largely unaffected. Accordingly, following evaluations were conducted on commodity hardware (one 2.5GHz CPU, 8 GB RAM) in a virtualised Linux environment:

- Biometric performance (DET curve) with the original and quantised feature vectors.
- Time elapsed (in ms) for the computations in the HE domain.

3.2 Results

When applying the original float-based templates and CKKS scheme, the distance computation between two encrypted feature vectors was around 5000ms. However, by applying a quantisation scheme, which enables the use of BFV, significant speed-up was achieved – the distance computation only taking 850ms. The time consumed by the encryption and

decryption operations was trivial in comparison to that of the distance computation, taking around 7ms and 2.5ms for CKKS and BFV, respectively. The space requirements for the generated keys are not excessive: <1MB for the public, secret, and relinearization keys (each) and around 10MB for the Galois keys. In figure 2, it can be seen that the biometric performance of the system was not degraded by the application of quantisation.




Fig. 2: DET curves for biometric identification with original and quantised features (with 95% CI)

4 Technical Considerations

By utilising HE, the security objectives of a biometric template protection system (see section 1) are achieved. The unlinkability across different databases can be guaranteed, insofar they use a different set of keys for the encryption. The irreversibility of the templates is bound to the encryption strength, which in the used library, depending on the chosen parameters, can be 128, 192, or 256 bits. Renewability is ensured, since new protected templates can always be generated by changing the encryption keys. Finally, as the template comparator in the encrypted domain is functionally identical (yields the same comparison scores) to that of the plaintext domain, the biometric performance is not impacted.

The speed of the current implementation may be prohibitive for larger deployments. It should, however, be noted that the experiments were carried out using an ordinary set-up, *i.e.* without powerful CPUs, parallelisation/distribution, *etc.* The hardware set-up notwithstanding, it would be beneficial to incorporate the concepts of computational workload reduction (see *e.g.* [KPB18]) into such systems in order to narrow down the search space for each identification transaction. However, one drawback of using HE is that it limits the flexibility in the implementation – for instance, as previously mentioned, due to batching the feature vector elements cannot be accessed individually. This limitation makes *e.g.* the incremental recognition schemes (which facilitate early acceptance/rejection of likely/unlikely candidates) infeasible. The incorporation of more complicated schemes, such as indexing and binning, could be a potentially interesting future research avenue. On the other hand, a 1-to-first search strategy could already be implemented by slight alterations to the communication between the database and the trusted third party described in section 2, however likely at the cost of at least some information exposure.

In the evaluation of biometric systems, a multitude of factors need to be considered. Some of the most important properties are the biometric performance, computational workload, as well as data security and privacy. Those goals typically counterbalance each other, and a biometric system operator is inevitably faced with trade-offs. In the case of HE in the biometric identification scenario described in this paper, currently the goals of high biometric performance, as well as data security and privacy can be fulfilled, while reducing or accelerating the computational workload could be pursued in future research to facilitate usage of such HE schemes in the practical, real-time applications.

5 Summary

In this paper, an architecture and an implementation thereof for a facial identification system in HE domain were presented and subsequently evaluated experimentally. The system fulfils the biometric template protection objectives defined in ISO/IEC IS 24745, namely unlinkability, irreversibility, renewability, and does not negatively affect the biometric performance. The paper also provided a discussion w.r.t. the technical considerations, challenges, and future work potential for utilisation of HE in the context of biometric face identification systems. While not yet practical for real-time applications, HE is definitely a promising avenue for future research and developments in this context.

Acknowledgements

This work was partially supported by the German Federal Ministry of Education and Research (BMBF), by the Hessen State Ministry for Higher Education, Research and the Arts (HMWK) within Center for Research in Security and Privacy (CRISP), and the LOEWE-3 BioBiDa Project (594/18-17). The authors acknowledge the financial support by the Federal Ministry of Education and Research of Germany in the framework of MEDIAN (FKZ 13N14798).

References

- [Ac18] Acar, A.; Aksu, H.; Uluagac, A. S.; Conti, M.: A Survey on Homomorphic Encryption Schemes: Theory and Implementation. *ACM Comput. Surv.*, 51(4):79:1–79:35, 2018.
- [Ag13] Aguilar-Melchor, C.; Fau, S.; Fontaine, C.; Gogniat, G.; Sirdey, R.: Recent Advances in Homomorphic Encryption: A Possible Future for Signal Processing in the Encrypted Domain. *IEEE Signal Processing Magazine*, 30(2):108–117, March 2013.
- [BGN05] Boneh, D.; Goh, E.-J.; Nissim, K.: Evaluating 2-DNF Formulas on Ciphertexts. In: *Second Theory of Cryptography Conference*. volume 3378 of *Lecture Notes in Computer Science*. Springer, pp. 325–341, 2005.
- [BGV14] Brakerski, Z.; Gentry, C.; Vaikuntanathan, V.: (Leveled) Fully Homomorphic Encryption without Bootstrapping. *TOCT*, 6(3):13:1–13:36, 2014.
- [Bo19] Boddeti, V. N.: Secure Face Matching Using Fully Homomorphic Encryption. In: *Intl. Conf. on Biometrics Theory, Applications and Systems (BTAS)*. IEEE, pp. 1–10, 2019.
- [BV11] Brakerski, Z.; Vaikuntanathan, V.: Fully Homomorphic Encryption from Ring-LWE and Security for Key Dependent Messages. In: *31st Annual Cryptology Conference*. volume 6841 of *Lecture Notes in Computer Science*. Springer, pp. 505–524, 2011.
- [Ca13] Campisi, P.: *Security and privacy in biometrics*, volume 24. Springer, 2013.
- [Ch17] Cheon, J. H.; Kim, A.; Kim, M.; Song, Y.: Homomorphic Encryption for Arithmetic of Approximate Numbers. In: *ASIACRYPT*. Springer, pp. 409–437, 2017.

- [Dr18] Drozdowski, P.; Struck, F.; Rathgeb, C.; Busch, C.: Benchmarking Binarisation Schemes for Deep Face Templates. In: Intl. Conf. on Image Processing (ICIP). IEEE, pp. 1–5, 2018.
- [EG85] El-Gamal, T.: A public key cryptosystem and a signature scheme based on discrete logarithms. *Trans. Information Theory*, 31(4):469–472, 1985.
- [Eu16] European Parliament: Regulation (EU) 2016/679. *Official Journal of the European Union*, L119:1–88, April 2016.
- [FV12] Fan, J.; Vercauteren, F.: Somewhat Practical Fully Homomorphic Encryption. *IACR Cryptology ePrint Archive*, 2012:144, 2012.
- [Ge09] Gentry, C.: A Fully Homomorphic Encryption Scheme. PhD thesis, 2009.
- [GH11] Gentry, C.; Halevi, S.: Implementing Gentry’s fully-homomorphic encryption scheme. In: Annual intl. conf. on the theory and applications of cryptographic techniques. Springer, pp. 129–148, 2011.
- [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.
- [IS11] ISO/IEC JTC 1/SC 27 IT Security techniques: . ISO/IEC 24745:2011. Information technology – Security techniques – Biometric information protection. International Organization for Standardization and International Electrotechnical Committee, June 2011.
- [ISR18] Imtiyazuddin, S.; Subba Rao, Y. V.; Rekha, N. R.: Faster Biometric Authentication System using Fan and Vercauteren Scheme. In: Intl. Conf. on Advances in Computing, Control and Communication Technology (IAC3T). pp. 48–53, September 2018.
- [KPB18] Kavati, I.; Prasad, M.; Bhagvati, C.: Search space reduction in biometric databases: A review. In: *Computer Vision: Concepts, Methodologies, Tools, and Applications*, pp. 1600–1626. IGI Global, 2018.
- [Mi19] Microsoft SEAL (release 3.2). <https://github.com/Microsoft/SEAL>, February 2019. Microsoft Research, Redmond, WA.
- [NJ15] Nandakumar, K.; Jain, A. K.: Biometric Template Protection: Bridging the performance gap between theory and practice. *IEEE Signal Processing Magazine*, 32(5):88–100, September 2015.
- [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.
- [PRC15] Patel, V. M.; Ratha, N. K.; Chellappa, R.: Cancelable biometrics: A review. *IEEE Signal Processing Magazine*, 32(5):54–65, 2015.
- [RSA78] Rivest, R. L.; Shamir, A.; Adleman, L. M.: A Method for Obtaining Digital Signatures and Public-Key Cryptosystems. *Commun. ACM*, 21(2):120–126, 1978.
- [RU11] Rathgeb, C.; Uhl, A.: A survey on biometric cryptosystems and cancelable biometrics. *EURASIP J. on Information Security*, 2011.
- [SKP15] Schroff, F.; Kalenichenko, D.; Philbin, J.: FaceNet: A unified embedding for face recognition and clustering. In: *Conf. on Computer Vision and Pattern Recognition (CVPR)*. IEEE, pp. 815–823, 2015.
- [Ya82] Yao, A. C.: Protocols for Secure Computations (Extended Abstract). In: 23rd Annual Symposium on Foundations of Computer Science. IEEE Computer Society, pp. 160–164, 1982.