Pattern Recognition 67 (2017) 149–163

Contents lists available at ScienceDirect

Pattern Recognition

journal homepage: www.elsevier.com/locate/patcog

Multi-biometric template protection based on Homomorphic Encryption

Marta Gomez-Barrero[a,*], Emanuele Maiorana[b], Javier Galbally[c], Patrizio Campisi[b], Julian Fierrez[d]

[a] da/sec - Biometrics and Internet Security Research Group, Hochschule Darmstadt, Germany
[b] Dipartimento di Ingegneria, Sezione di Elettronica Applicata, Universita degli Studi Roma Tre, Italy
[c] European Commission - Joint Research Centre, Italy
[d] ATVS - Biometric Recognition Group, EPS, Universidad Autonoma de Madrid, Spain

ARTICLE INFO

Article history:
Received 20 June 2016
Revised 14 October 2016
Accepted 15 January 2017
Available online 21 January 2017

Keywords:
Biometrics
Template protection
Multi-biometrics
Homomorphic Encryption
Unlinkability
Privacy

ABSTRACT

In spite of the advantages of biometrics as an identity verification technology, some concerns have been raised due to the high sensitivity of biometric data: any information leakage poses a severe privacy threat. To solve those issues only protected templates should be stored or exchanged for recognition purposes. In order to improve the performance and achieve more secure and privacy-preserving systems, we propose a general framework for multi-biometric template protection based on homomorphic probabilistic encryption, where only encrypted data is handled. Three fusion levels are thoroughly analysed, showing that all requirements described in the ISO/IEC 24745 standard on biometric data protection are met with no accuracy degradation. Furthermore, even if all the process is carried out in the encrypted domain, no encryptions are necessary during verification, thereby allowing an efficient verification which can be deployed for real-time applications. Finally, experiments are carried out on a reproducible research framework. The results obtained show high accuracy rates, reaching EERs as low as 0.12%, and requiring protected templates comprising 200 KB.

1. Introduction

Over the last decades, biometric recognition has been developed as a reliable alternative to traditional authentication systems based on PINs or passwords [1]. Among other advantages, biometric characteristics (e.g., face, fingerprint or signature) cannot be lost or forgotten. On the other hand, biometric information is very sensitive and some concerns have been raised regarding the privacy of the subjects - it has already been proved that samples can be recovered from unprotected templates [2–5] and subsequently used to impersonate genuine subjects [6]. In fact, biometric data is considered sensitive data in European Union (EU) General Data Protection Regulation 2016/679 [7], which means that the use of these data is subjected to the right of privacy preservation. As a consequence, templates must be protected, in order to prevent any leakage of this highly sensitive information.

Such privacy protection is granted by *Biometric Template Protection* (BTP) schemes [8,9]. In these systems, a protected template **C**, which should not reveal any biometric information, is extracted from the biometric sample **M**, possibly using a secret key, and stored in the database. More specifically, the ISO/IEC 24745 standard on biometric information protection [10] establishes two main requirements for such templates:

- *Irreversibility*: "biometric data shall be processed by irreversible transforms before storage". Therefore, given a protected template **C**, no biometric information should be learned from it. It should also not be feasible to reconstruct a biometric sample $\mathbf{M}'$, which is positively matched to the original sample **M** by the biometric system.
- *Unlinkability*: " the stored biometric references should not be linkable across applications or databases". In other words, given two templates, $\mathbf{C}_1$ and $\mathbf{C}_2$, extracted from the same biometric sample **M** and protected with different secret keys, it should not be feasible to decide whether they belong to the same subject.

Additionally, due to the fact that biometric characteristics cannot be replaced, *renewability* is also desired (i.e., if one template is

* Corresponding author.
E-mail addresses: marta.gomez-barrero@h-da.de (M. Gomez-Barrero), maiorana@uniroma3.it (E. Maiorana), javier.galbally@jrc.ec.europa.eu (J. Galbally), campisi@uniroma3.it (P. Campisi), julian.fierrez@uam.es (J. Fierrez).

http://dx.doi.org/10.1016/j.patcog.2017.01.024
0031-3203/

lost or stolen, a new one, not matching the old template, should be issued). At the same time, other properties such as verification accuracy, speed and storage requirements should be maintained compared to the same system using unprotected data [11].

Most existing BTP schemes [12,13], commonly categorized as cancelable biometrics and cryptobiometrics, show at least one of the following two drawbacks: *i*) performance degradation with respect to unprotected systems; *ii*) they require *Auxiliary Data* (AD) for verification purposes [8]. Attacks on this AD can disclose sensible information, which compromises both the security of the system and the privacy of the subject [14,15].

As an alternative to those approaches using AD, Homomorphic Encryption schemes allow for computations to be performed on ciphertexts, with no additional AD, and which generate encrypted results which decrypt to plaintexts that match the result of the operations carried out on the original plaintext [16]. Therefore, combining such an encryption approach with biometric verification systems would meet the aforementioned requirements while preserving verification performance [17]. Since practical implementations of Fully Homomorphic Encryption (FHE) schemes still remain a big challenge [18], somewhat Homomorphic Encryption (HE) schemes, which only allow a limited subset of operations in the encrypted domain, are nowadays being introduced into many applications based on signal processing [18,19], and, particularly, biometrics [20–22].

However, as HE only allows a restricted set of operations in the encrypted domain, it is possible that some of the most advanced and accurate state-of-the-art systems, such as GMMs or SVMs, are difficult to integrate in the proposed framework while keeping verification time low enough for real time applications [23]. This limitation can be overcome by introducing *multi-biometric* template protection schemes (MBTP) [24], since the combination of different biometric characteristics generally leads to higher accuracy [25]. As defined in the ISO/IEC TR 24722 on multimodal and other multibiometric fusion [26], fusion can be carried out at three different levels [25], namely:

- *Feature level fusion*: a single template of higher dimensionality is generated from the individual templates extracted from each characteristic, hence comprising more discriminative information than each single template.
- *Score level fusion*: each unimodal system returns an individual similarity score, which are normalized to a common range and combined in order to obtain a more accurate system [27].
- *Decision level fusion*: each unimodal system returns an individual accept/reject decision, which are fused in order to increase the accuracy of the system.

Even though extensive research has been carried out on the fields of multi-biometric recognition [25] and *unimodal* biometric template protection [8], several issues remain unsolved in the development of *multi-biometric* template protection schemes [28]. Two of the most significant challenges are: *i*) the development of a generic framework for multi-biometric template protection, and *ii*) the difficulty to obtain protected templates from non pre-aligned samples, without requiring AD (and hence avoiding potential information leakage).

In the present article, we propose and implement a general template protection framework for multi-biometric systems based on Homomorphic Encryption, analysing the advantages and disadvantages of each fusion level. Two different distance measures are implemented for this particular encryption scheme and compared in terms of: verification performance, irreversibility, unlinkability and computational complexity. To the best of our knowledge, this is the first MBTP scheme based on HE.

Additionally, the proposed MBTP scheme is defined within an overall security model (see Section 3) that describes the entities involved in the recognition process and their expected behaviour (i.e., honest or malicious). This way the reader can have a clearer picture of the threats being taken into account and what are the risks against which the protection approach is effective.

Finally, in order to assess the soundness of the proposed approach, we evaluate a particular case study based on the fusion of on-line signature and fingerprints. We have chosen this particular fusion due to its possible applications: while on-line signature is one of the most widely accepted biometric characteristics, fingerprints offer a very high accuracy and are being deployed in most smartphones. Furthermore, the experimental evaluation is carried out on a reproducible research framework: the real and publicly available multimodal BiosecurID database [29] and an open source implementation of the Paillier cryptosystem.[1] Results are reported for the protected and unprotected domains following a common protocol.

In summary, most proposed BTP schemes report a degradation in verification performance. Furthermore, should a protected template be stolen, there is no way to recover the original biometric sample, thus requiring the re-acquisition of biometric samples in order to re-generate a biometric protected database. In addition, for a subset of methods, AD has to be stored together with the protected templates, being hence more vulnerable to attacks such as the ones proposed in [14,15]. The method proposed in the present article deals with those drawbacks by using Homomorphic Encryption to encrypt biometric templates and the computations carried out at verification time. Furthermore, verification accuracy is enhanced thanks to a multi-biometric approach.

2. Related works

Biometric template protection schemes [12,13] have been traditionally divided into *cancelable biometrics* and *cryptobiometrics*. Homomorphic Encryption and Garbled circuits do not belong to these categories. A third class, namely *biometrics in the encrypted domain*, can therefore be added. The general classification of BTP schemes considered in the present article is depicted in Fig. 1.

Sections 2.1 and 2.2 focus on the review of protection schemes proposed for *multi-biometric* systems. The references in Fig. 1 correspond to the works described in those sections. However, as depicted in Fig. 1, not all the methods developed for unimodal systems have been translated to the multi-biometric case. For an exhaustive review of works dealing with *cancelable biometrics* and *cryptobiometrics* template protection in *unimodal* systems, the reader is referred to [8].

On the other hand, as no schemes have been presented so far for multimodal systems in the *biometrics in the encrypted domain* class, Section 2.3 reviews the existing protection methods in this category for *unimodal* systems. A summary of the main characteristics of each method is included in Table 1. Even if some properties had been proven in the corresponding articles, if any attack has later shown that the system was not irreversible or not unlinkable, we have written " No". Therefore, as it may be observed, none of the methods grant the aforementioned requirements for BTP schemes (i.e., irreversibility, unlinkability and no performance degradation) at the same time.

2.1. Cancelable multi-biometrics

Cancelable biometrics refer to schemes in which biometric data is obscured with an irreversible transformation and verification is carried out in the transformed domain. As shown in Fig. 1, there are two main types of schemes: *i*) irreversible transformations of

[1] http://www.csee.umbc.edu/~kunliu1/research/Paillier.html

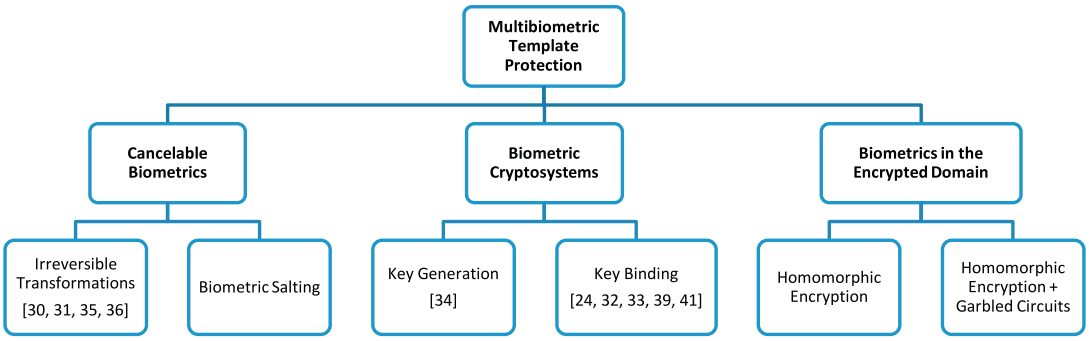

Fig. 1. Multi-biometric template protection taxonomy.

Table 1
Summary of key multi-biometric template protection schemes. " –" stands for properties not evaluated in the corresponding articles.

Technique	Ref.	Accuracy	Accuracy drop	Irreversibility	Unlinkability	Complexity
Cryptobiometrics: Fuzzy Vault & Commit.	[24]	68% GAR 75% GAR	–	No	No	Low
Cancelable: Random projection	[30]	89% GAR 0% FMR	–	Yes	–	Low
Cancelable: BioHashing, Interp. BioConvolving	[31]	∼ 95% GAR	∼ 5%	Yes	–	Low
Cryptobiometrics: Fuzzy vault	[32]	1.8% FMR 0.01% FNMR	–	No	No	Low
Cryptobiometrics: Cascade	[33]	0% EER	–	Yes	–	Low
Cryptobiometrics: Quantization	[34]	0.92% FMR 0.0002% FNMR	–	Yes	–	Low
Cancelable: Irrev. Transformation	[35]	1.9% FMR 0.01% FNMR	0%	Yes	No	Low
Cancelable: Irrev. Transformation	[36]	6% EER	–	Yes	–	Medium
Cryptobiometrics: Fuzzy Comm.	[37]	∼ 3% EER	–	No	No	Low
Cryptobiometrics: Cascade	[38]	–	–	No	No	Low

the biometric data or unprotected templates, and *ii*) biometric salting, in which AD is blended with biometric data to derive a distorted version of the biometric template. However, MBTP have only been proposed for the former. A fusion of face and ear samples using random projections and a transformation-based feature extraction, reducing dimensionality with PCA and clustering the features, is proposed in [30]. Applying different cancelable transformations, voice and iris are fused in [31]. The spiral and the continuous components belonging to two different fingerprints from the same subject are mixed in one cancelable template in [36]. More recently, face and iris features are protected with Bloom filters and fused into a single template in [35].

All of the cited approaches, except the scheme proposed in [35], lead to a degradation in verification performance. Furthermore, should a protected template be stolen, there is no way to recover the original biometric sample or the unprotected template in order to re-encode it with a new key. As a consequence, in order to re-generate the biometric database, biometric samples need to be re-acquired, with the additional nuisance this fact could pose to the subjects.

2.2. Multi-biometric cryptosystems

On the other hand, in cryptobiometric systems a key is either bound (key binding schemes in Fig. 1) or extracted (key generation schemes in Fig. 1) from biometric data. In this context, most systems rely on the fuzzy vault [39] and the fuzzy commitment [40] schemes, which are classified as key binding approaches.

A fuzzy vault scheme is proposed in [32], where a single multi-biometric template is derived from fingerprint and iris features. On the other hand, a fuzzy commitment scheme for the fusion of two different feature extraction algorithms is applied to 3D face data in a single sensor scenario in [37], which is therefore not applicable to the more general fusion of different biometric modalities. Additionally, fuzzy schemes can be applied to secure sketches, that is, secure representations of biometric templates in which helper data is used to recover the original biometric template and matching is reduced to an error correction [41]. In [24], a single secure sketch is generated from multiple and heterogeneous templates, based on the concatenation of the individual sketches. Practical implementations for fuzzy vault and fuzzy commitment schemes are then proposed for the fusion of iris, fingerprint and face.

Those methods share a common drawback: AD has to be stored together with the protected templates, being hence more vulnerable to attacks such as the ones proposed in [14,15].

On the other hand, regarding key generation systems, quantization schemes are applied to face and pre-aligned fingerprint samples in [34].

In opposition to those previous schemes, a modular approach for the design of multi-biometric cryptosystems is proposed in [33]: a secure sketch is extracted from each biometric template, and used in a sequential manner to secure successive templates. In [38], a more general approach is presented, where multiple secrets are similarly used in a cascade fashion within the secure sketch framework. In this last case, no evaluation of the verification performance is provided. These approaches, classified as key binding

schemes, have the advantage of an easy escalation to more biometric samples, while the main limitation is that the overall security is bounded by the security of the outermost layer.

As in the case of cancelable biometrics, cryptobiometric systems usually present a performance degradation with respect to the systems relying on unprotected data.

2.3. Biometrics in the encrypted domain

As an alternative to the aforementioned approaches, where either the privacy of the subject is not fully protected or verification performance degrades, secure multiparty computation and homomorphic cryptosystems can be used in order to carry out biometric recognition in the encrypted domain while obtaining results fully comparable to those yielded by plain data [18,42]. In particular, current approaches to *biometrics in the encrypted domain* [22] are based on Garbled Circuits (GC) [43] and Homomorphic Encryption (HE) [16,18] (see Fig. 1).

Since efficient implementations of HE schemes are very recent [44], only a few *unimodal* biometric systems based on this protection technology have been proposed so far. In [21], the authors present a new fingerprint verification system based on the FingerCode fixed-length representation of fingerprints [45] and HE. Results show that verification performance is preserved. However, the database stored in the server is not encrypted and results are reported on a small database comprising data of only 51 subjects. An improved version of that approach is suggested in [46], where a more compact implementation using quantization is proposed at a small cost in terms of verification accuracy.

In [47], Eigenface based templates are protected with HE. Then, a more efficient approach is presented in [48] using GCs for the threshold comparison. Furthermore, the SCiFI project [49] proposes a biometric identification algorithm specifically designed for a more efficient usage in secure computation, based on fixed-length templates with a constant Hamming weight. In contrast to the Eigenface based approaches, only the matching process, but not the template construction, is secured.

In [50], a secure iris BTP based on a combination of HE and GCs is proposed, handling encrypted iriscodes. In order to deal with the computation of Hamming Distances in the encrypted domain (divisions are not supported), the division and comparison with a verification threshold is reduced to a inequality comparison carried out with GCs.

Furthermore, even though biometric algorithms that achieve better detection rates are known in the literature, these schemes are much more complex than the representations used in the aforementioned articles [23], and thus more difficult to implement in the encrypted domain due to the limitation in the number of possible operations that can be performed. In order to compensate for such loss on verification accuracy, the fusion of several biometric characteristics in the encrypted domain is proposed.

3. Security model

We will use the following notation in the subsequent sections:

- $\mathbf{T}_p = \{p_1, \ldots, p_f, \ldots, p_F\}$: unprotected templates, comprising F features p_f.
- $S_{dist} = d_{dist}(\mathbf{T}_p, \mathbf{T}_r)$: similarity score between two templates $\mathbf{T}_p$ and $\mathbf{T}_r$, where d_{dist} is a particular distance function: *euc* stands for Euclidean and *cos* for cosine (see Section 4.2).
- m and m^*: plain message and its corresponding ciphertext.
- $m^* = E_{pk}(m, s)$, where E denotes the encryption function, s a random number and pk the public key.
- $m = D_{sk}(m^*)$, where D denotes the decryption function and sk the private key.

In the present section we describe the general security model considered in this work, including all the assumptions made regarding the expected behaviour (honest or malicious) of each of the entities involved in the biometric recognition process. This way the reader can have a more general perspective of how different threats have been taken into account and how the proposed MBTP deals with several privacy and/or security risks.

First, a general diagram of the unencrypted biometric verification system is depicted in Fig. 2 (left), where two entities are involved:

- A client, which will acquire the probe biometric sample, extract the features and encode them in the template $\mathbf{T}_p$, generate the similarity score as the distance between $\mathbf{T}_p$ and the reference template $\mathbf{T}_r$, $S = d(\mathbf{T}_p, \mathbf{T}_r)$ (see Eqs. (6) and (9)), and compute the final genuine/impostor verification decision $D = (S > \delta)$, where δ is the pre-defined verification threshold.
- A server, which will hold the database with the reference templates $\mathbf{T}_r$ and send them to the client during verification.

In order to increase the privacy of the subject, the server must process the client's biometric data without extracting any information from it, and at the same time, the server must protect the information stored in the database [17]. To that end, a different security model is used in the protected system (see Fig. 2, right) where all the data, either stored or shared between client and server in the verification process, should be encrypted. Therefore, the new entities and roles are the following:

- The client acquires the probe biometric sample, extracts the template $\mathbf{T}_p$ and generates the encrypted score $E(S)$ (see Eqs. (7) and (10)), sending it to the server.
- The DB server holds the database comprising only encrypted templates ($E(\mathbf{T}_r)$) and sends the encrypted reference template $E(\mathbf{T}_r)$ to the client during verification.
- The authentication server holds the key pair (sk, pk) and computes the final genuine/impostor decision D.

With respect to the model proposed in other biometric template protection approaches [51], in the present scheme the sensor and matcher have been integrated into a single entity: the client.

Therefore, the requirements on the data flow described in [51] to fulfil the aforementioned irreversibility and unlinkability criteria have been adapted to the present model as:

- The authentication server should not learn $\mathbf{T}_r$ or $\mathbf{T}_p$.
- The database server should not learn $\mathbf{T}_p$ or trace subjects.
- The client should not learn $\mathbf{T}_r$.

To fulfil those requirements we assume that:

- According to the honest-but-curious adversary model [52], all parts involved follow the protocols honestly. As a consequence, we may assume that the scores computed by the client are correct.
- An adversary may have access to one of the servers, but the authentication and DB servers will not collude.

This way, since the client does not know sk, it cannot decrypt the reference template $E(\mathbf{T}_r)$ or the similarity score $E(S)$. The comparator can hence be moved to the server and S cannot be used to carry out hill-climbing attacks that need access to the score in order to be performed [53]. Additionally, as it will be shown in Section 4.2, the probe template $\mathbf{T}_p$ does not need to be encrypted, since it never leaves the client (and we are assuming an honest behaviour from the client). As a consequence there is no leak of biometric information in the communication channel.

In a similar manner, the authentication server does not have access to either the probe template $\mathbf{T}_p$ or the encrypted reference $E(\mathbf{T}_r)$. This way, it cannot learn any biometric data.

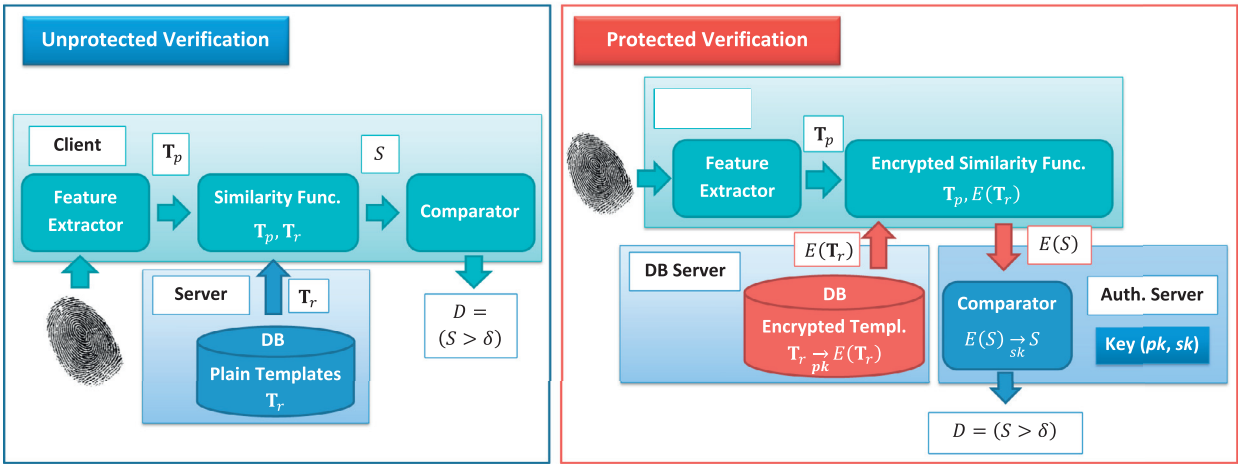

Fig. 2. Unprotected vs protected biometric verification. In the unprotected scenario (left), a probe biometric sample is acquired and its features extracted ($\mathbf{T}_r$). The final output is the binary decision $D = (S > \delta)$, where S is computed as the distance with respect to the stored unprotected reference ($\mathbf{T}_p$): $S = d(\mathbf{T}_r, \mathbf{T}_p)$, according to Eqs. (6) and (9) (depending on the distance measure considered). In the protected scenario (right), all the encrypted data or information flow is depicted in red: $E(\mathbf{T}_p)$ and $E(S)$. Analogously, the similarity function in the encrypted domain is defined in the text by Eqs. (7) and (10). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

4. Proposed system

Following the security model proposed in the previous section, the Paillier homomorphic probabilistic encryption scheme [44] will be used to encrypt the data (more details in Section 4.1). The implementation of two different distance functions (i.e., Euclidean and Cosine) which reached an optimum accuracy compared to other distances for this particular encryption scheme is described in Section 4.2. Finally, in Section 4.3 we propose a general framework in the encrypted domain for each fusion level defined in the ISO/IEC TR 24722 on multimodal and other multi-biometric fusion [26].

4.1. Homomorphic Encryption

Among the HE schemes proposed in the literature, the Paillier homomorphic probabilistic encryption scheme [44] is used in the present work, due to its advantages over other cryptosystems for this particular application. For instance, the Goldwasser–Micali cryptsosystem [54] is homomorphic for the XOR operation but not for the addition, as needed for the distance functions computations. Other widely used schemes with aditively homomorphic properties include ElGamal crypstosystem [55] and its variants. However, as shown in [44], decryption is faster for the Paillier cryptosystem (and as we will see in Section 4.3 one decryption is carried out during verification). Furthermore, Paillier also offers homomorphic multiplication of plaintexts, needed for the distance functions considered.

Paillier cryptosystem is based on the decisional composite residuosity assumption (DCRA): given a composite n and an integer z, it is hard to decide whether z is an n-residue modulo n^2. In other words, it is hard to decided whether there exists y such that $z = y^n \mod n^2$.

As any other public key encryption scheme, two separate keys are required: *i*) a public key pk for encryption, and *ii*) a private or secret key sk for decryption. In the Paillier cryptosystem, the public key is defined as $pk = (n, g)$, where $n = p \cdot q$ with p and q two large prime numbers such that $\gcd(pq, (p-1)(q-1)) = 1$, and $g \in \mathbb{Z}^*_{n^2}$ (i.e., the set of integers modulo n^2 which possess a multiplicative inverse). On the other hand, the secret key is defined as $sk = (\lambda, \mu)$, where $\lambda = \text{lcm}(p-1, q-1)$ and $\mu = (g^\lambda \mod n^2)^{-1} \mod n$.

Given a message $m \in \mathbb{Z}_n$, its encryption is denoted as $m^* = E_{pk}(m, s) \in \mathbb{Z}^*_{n^2}$, and computed as follows:

$$E_{pk}(m, s) = g^m \cdot s^n \mod n^2 \tag{1}$$

where $s \in \mathbb{Z}^*_n$ is a random number providing the probabilistic nature of the cryptosystem. This property is necessary to grant semantic security against chosen-plaintext attacks [54]. In particular, different ciphertexts are obtained when the same plaintext is encrypted several times using the same public key: $E_{pk}(m, s_1) \neq E_{pk}(m, s_2)$. This randomness provides the required unlinkability to the protected templates: even if the exact same unprotected features are extracted from a particular biometric sample, the encrypted templates would be different.

It is shown in [44] that E is a one-way function (i.e., irreversible) if and only if the decisional composite residuosity assumption holds. Therefore, a computationally-bound attacker in possession of an encrypted message m^* (a protected biometric template) and the public key pk would not be able to extract any information about the plaintext m (biometric information). He could only do so if he obtained the secret key sk and decrypted the ciphertext $m^* = E_{pk}(m, s)$ as follows

$$m = D_{sk}(m^*) = L\left((m^*)^\lambda \mod n^2\right) \cdot \mu \mod n \tag{2}$$

where $L(t) = (t-1)/n$.

The main advantage of HE schemes with respect to other cryptosystems is the fact that some operations can be carried out in the encrypted domain, yielding ciphertexts whose corresponding plaintexts are the same we would obtain performing the operations over the plaintexts. In particular, the Paillier cryptosystem fulfils two properties which will be used in the present scheme. On the one hand, the product of two ciphertexts, $m^*_1 = E_{pk}(m_1, s_1)$ and $m^*_2 = E_{pk}(m_2, s_2)$, will decrypt to the sum of their corresponding plaintexts:

$$D_{sk}\left(m^*_1 \cdot m^*_2 \mod n^2\right) = m_1 + m_2 \mod n \tag{3}$$

On the other hand, an encrypted plaintext, $m^*_1 = E_{pk}(m_1, s_1)$, raised to a constant l, will decrypt to the product of the plaintext and the constant:

$$D_{sk}\left((m^*_1)^l \mod n^2\right) = m_1 \cdot l \mod n \tag{4}$$

As a consequence, while an unlimited number of summations can be carried out in the encrypted domain, only a limited number of products can be computed - as it is shown in Eq. (4), one of the factors should be a plaintext. This fact poses a severe challenge for the implementation of many similarity measures.

4.2. Encrypted distance computation

In order to compare two biometric samples $\mathbf{T}_p$ and $\mathbf{T}_r$, two different distances will be considered. In the next subsections, we describe the implementation of each distance function within the Pailler cryptosystem. Since square roots and divisions are not straightforward to implement in the encrypted domain, we will use the square of the Euclidean distance. Furthermore, to avoid overcomplicated notation, the encrypted values $E_{pk}(m, s)$ will be simply denoted as $E(m)$, even though the random number s and the public key pk are needed for the encryption computation (see Eq. (1)). This way, for each particular distance:

- $\mathbf{T}_p = \{p_1, \ldots, p_F\}$ denotes the probe biometric sample.
- $E(\mathbf{T}_r)_{dist}$ denotes the encrypted reference template, for each distance measure. As will be explained in the following subsections, and defined in Eqs. (8) and (11), the encrypted template $E(\mathbf{T}_r)_{dist}$ is different for each distance. The reader should thus be aware that $E(\mathbf{T}_r) \neq \{E(r_1), \ldots, E(r_F)\}$. This is due to the impossibility to carry out some operations, such as division or square roots, in the encrypted domain. One of the contributions of the paper is defining $E(\mathbf{T}_r)_{dist}$ for each distance measure, so that the score can be directly computed in the encrypted domain.
- $E(S_{dist})$ denotes the encrypted similarity score, computed between $\mathbf{T}_p$ and $E(\mathbf{T}_r)_{dist}$ as defined in Eqs. (7) and (10). Following $E(\mathbf{T}_r)_{dist}$, a contribution of the paper is defining, for each of the two considered distance measures, the function $E(S_{dist})$ that takes as input $\mathbf{T}_p$ and $E(\mathbf{T}_r)_{dist}$, and outputs the encrypted score with no decryptions involved.

We need to take into account two limitations of the Paillier cryptosystem for the computation of $E(S_{dist})$: *i*) we can carry out a limited set of operations in the encrypted domain (unlimited sums but a limited number of products), and *ii*) we can only work with integers. Additionally, all features should be in the same value range in order to carry out the fusion of several modalities in the multi-biometrics system. To that end, we will consider a two-step approach. First, the real-valued extracted features will be normalized to the interval [0, 1]. Then, we will transform those normalized real-valued features to integer values in a bigger range, in our experiments $[0, 10^3]$, in order to retain as much information as possible:

$$X \rightarrow \text{round}\left(10^3 \; X\right) \tag{5}$$

4.2.1. Encrypted Euclidean distance

Given two F-dimensional templates $\mathbf{T}_p$ and $\mathbf{T}_r$, in the unprotected domain the score $S_{euc} = d^2_{euc}(\mathbf{T}_p, \mathbf{T}_r)$, can be efficiently computed as

$$S_{euc} = \sum_{f=1}^{F} p_f^2 + r_f^2 - 2 p_f r_f \tag{6}$$

Then, using Eqs. (3) and (4), the encrypted score can be directly computed in the encrypted domain without performing any encryptions in the client (Fig. 2 right) as

$$\begin{aligned} E(S_{euc}) &= \textstyle\prod_{f=1}^{F} E(1)^{p_f^2} \cdot E\left(r_f^2\right) \cdot E\left(r_f\right)^{-2p_f} \\ &= \textstyle\prod_{f=1}^{F} (1^*)^{p_f^2} \cdot euc2_f^* \cdot \left(euc1_f^*\right)^{-2p_f} \end{aligned} \tag{7}$$

The subject's reference template stored in the encrypted database is thus defined by the following ciphertexts:

$$E(\mathbf{T}_r)_{euc} = \{1^*\} \cup \left\{euc2_f^*, euc1_f^*\right\}_{f=1}^{F} \tag{8}$$

where $euc1_f^* = E(r_f)$ and $euc2_f^* = E(r_f^2)$. As a consequence, all cyphertexts involved in Eq. (7) are sent by the server, and products and exponentiations locally computed on the client.

It should be noted that, given the probabilistic nature of the Paillier cryptosystem, $E(1)$ can be computed and stored separately for each subject at enrolment time, leading to different encrypted values and thereby increasing the security and privacy of the subject. Furthermore, we will avoid the computational overhead caused by the encryption at verification time, at the cost of slightly increasing the storage requirements.

4.2.2. Encrypted cosine similarity

The cosine similarity between two F-dimensional vectors $\mathbf{T}_p$ and $\mathbf{T}_r$ is defined in the unencrypted domain (Fig. 2 left) as

$$d_{cos}(\mathbf{T}_p, \mathbf{T}_r) = \frac{\mathbf{T}_p \cdot \mathbf{T}_r}{\|\mathbf{T}_p\| \cdot \|\mathbf{T}_r\|} = \sum_{f=1}^{F} \frac{p_f \cdot p_f}{\|\mathbf{T}_p\| \cdot \|\mathbf{T}_r\|} \tag{9}$$

As proposed in [56], since $d_{\cos}(\mathbf{T}_p, \mathbf{T}_r)$ is a positive number in the range [0, 1], in order to have a bigger range of values that allows a comparison among integers with no significant information loss, we can compute the final similarity as $S_{\cos} = 10^{12} d_{\cos}(\mathbf{T}_p, \mathbf{T}_r)$, which can be directly computed in the encrypted domain (Fig. 2 right) as

$$E(S_{cos}) = \prod_{f=1}^{F} E\left(\frac{10^6 r_f}{\|\mathbf{T}_r\|}\right)^{10^6 p_f / \|\mathbf{T}_p\|} = \prod_{f=1}^{F} \left(cos_f^*\right)^{10^6 p_f / \|\mathbf{T}_p\|} \tag{10}$$

The subject's reference template stored in the encrypted database is therefore defined as

$$E(\mathbf{T}_r)_{cos} = \left\{cos_f^*\right\}_{f=1}^{F} \tag{11}$$

where the ciphertexts $cos_f^* = E(\frac{10^6 r_f}{\|\mathbf{T}_r\|})$. Therefore, all cyphertexts involved in Eq. (10) are sent by the server, and products and exponentiations locally computed on the client.

4.3. Encrypted multi-biometrics

This section builds upon the two encrypted distance measures described in Section 4.2 to present a new HE-based general multi-biometric template protection framework for each fusion level (i.e., feature, score and decision level). In order to avoid overcomplicated notation and with no loss of generality, we will stick to the case of fusing two biometric characteristics. However, it should be noted that the present framework can be applied to the fusion of any number of modalities.

4.3.1. Feature level fusion (Fig. 3)

At this level, a single protected template comprising all the features, related to both characteristics, is stored in the database and used at verification time. Therefore, all features are concatenated in a single encrypted template, and a single verification encrypted score $E(S)$ is computed. Identity verification is thus carried out in six steps, as shown in Fig. 3:

0. During enrolment, the reference biometric templates are encrypted using the server public key pk. The encrypted templates $E(\mathbf{T}_r^{fused})$ (see Eqs. (8) and (11)) are stored in the database.
1. Biometric samples are acquired and a single template, $\mathbf{T}_p^{fused}$, is extracted on the client.
2. The server sends the encrypted enrolled template to the client, $E(\mathbf{T}_r^{fused})$.
3. The client computes the similarity score in the encrypted domain: $E(S)$, according to Eqs. (7) and (10) (depending on the distance measure selected as encrypted similarity function).
4. The encrypted score $E(S)$ is sent to the server.

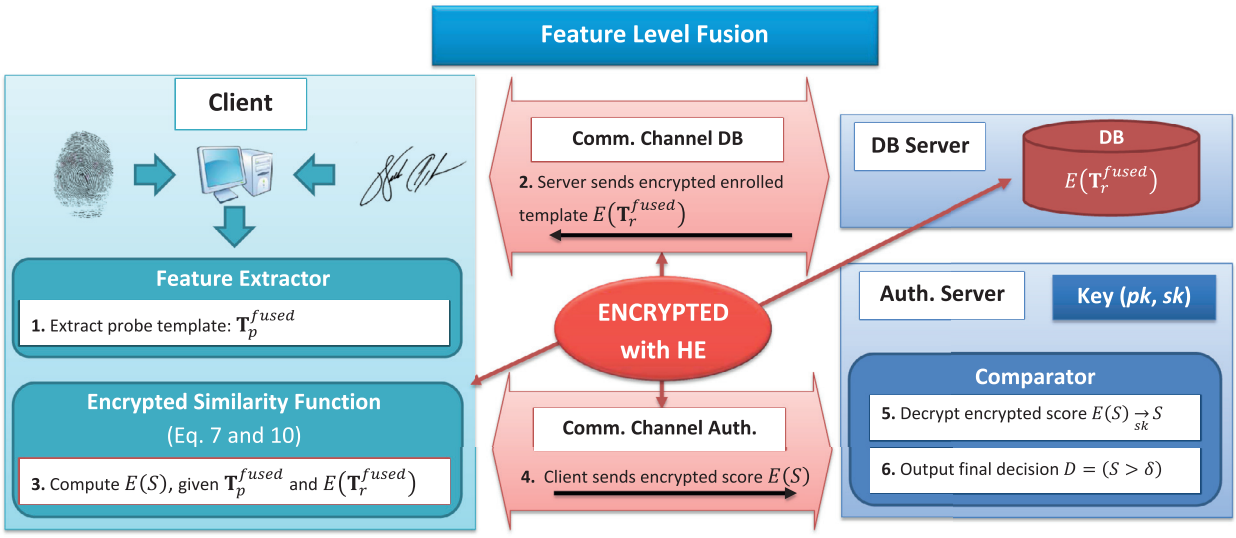

Fig. 3. General diagram of the proposed scheme: feature level fusion.

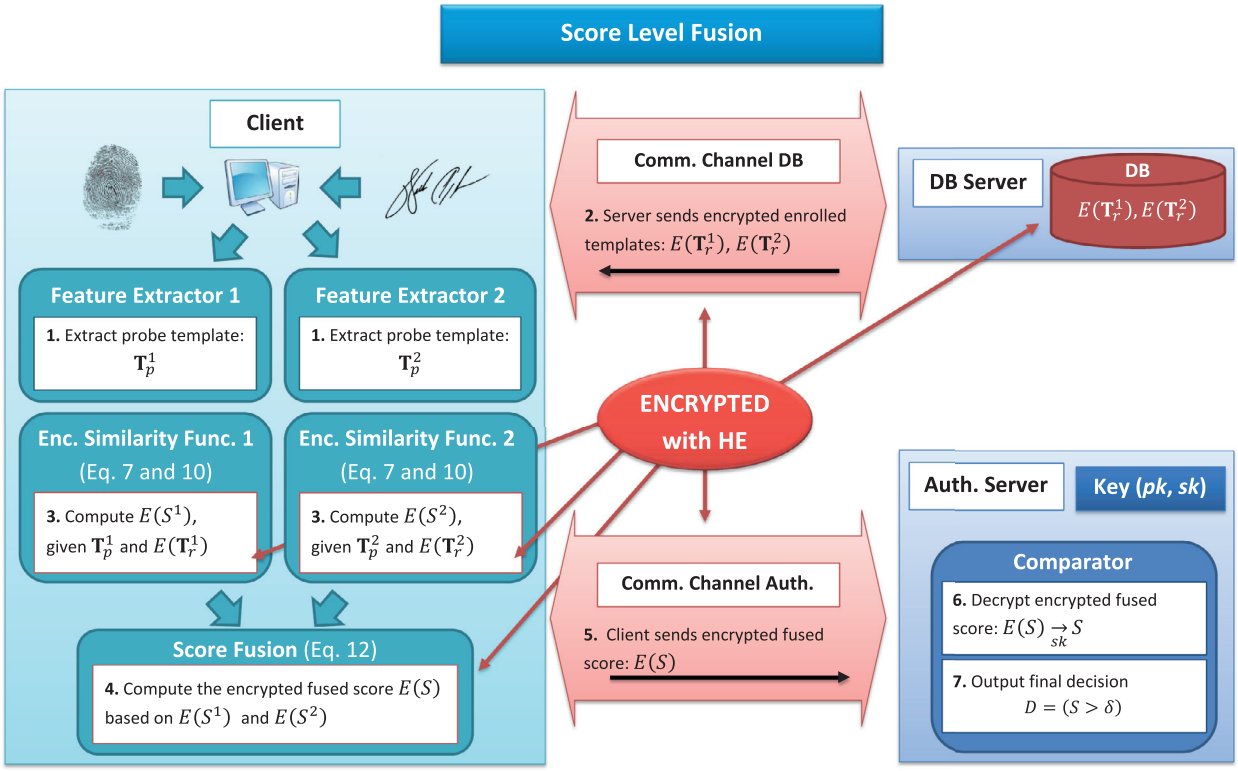

Fig. 4. General diagram of the proposed scheme: score level fusion.

5. The server decrypts the score using the secret key, *sk*.
6. Finally, the server compares the score to the verification threshold δ and outputs a genuine/impostor verification decision.

4.3.2. Score level fusion (Fig. 4)

In this approach, each biometric characteristic will be processed separately, generating two individual probe templates: $\mathbf{T}_p^1$ and $\mathbf{T}_p^2$. Similarly, the server stores and sends $E(\mathbf{T}_r^1)$ and $E(\mathbf{T}_r^2)$. The client matches them independently to $\mathbf{T}_p^1$ and $\mathbf{T}_p^2$ according to Eqs. (7) and (10), depending on the distance measure considered, producing two individual encrypted scores $E(S^1)$ and $E(S^2)$.

In order to normalise the individual scores prior to the fusion, several approaches are proposed in [57]. However, it is not possible to implement most of them in the encrypted domain without increasing the computational load due to the restriction in the type of operations that can be performed. We therefore propose a different and simpler approach, that achieves the same performance as the min-max rule proposed for the unprotected domain in [57]. Since all the scores are computed with the same distance measure and all the features are normalised to $[0, 10^3]$, for each particular distance the range of variation of the scores will depend on the dimensionality of the templates. Therefore, assuming that the dimensionality of the feature vector extracted from the first char-

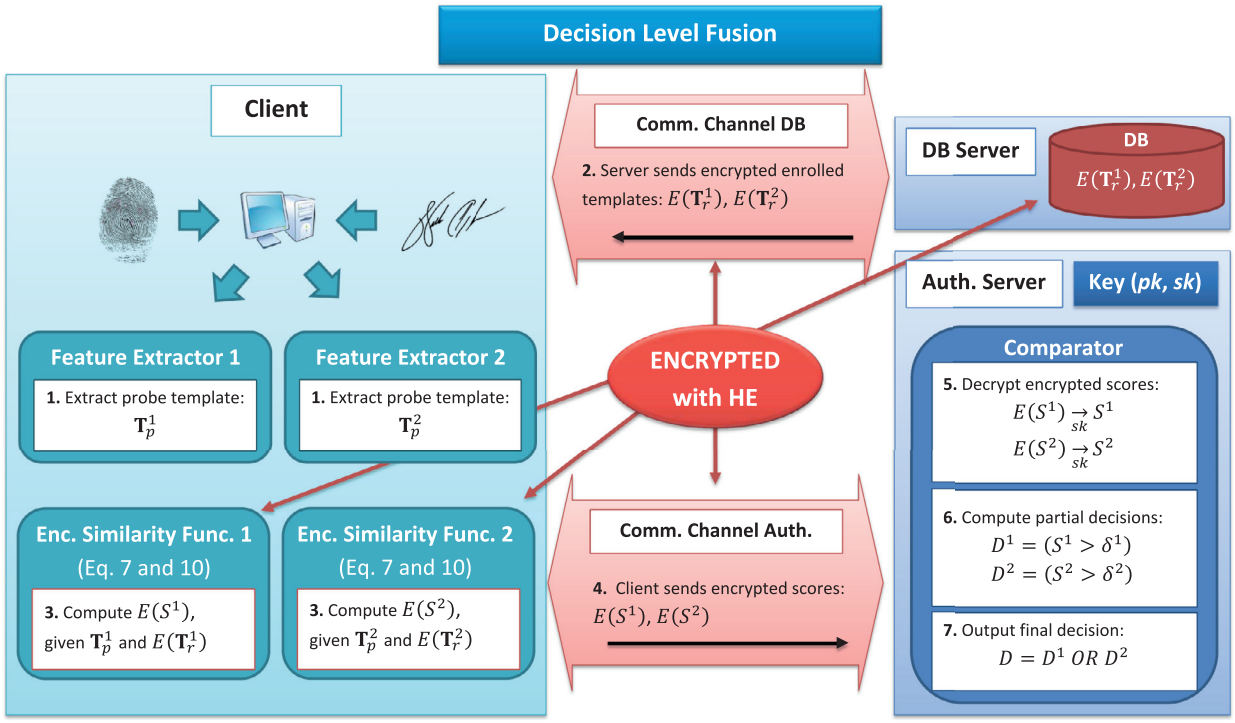

Fig. 5. General diagram of the proposed scheme: decision level fusion.

acteristic is higher ($F_1 > F_2$), we can perform the following normalisation, which in turn can be easily computed in the encrypted domain:

$$S'^2 = \beta S^2 \Rightarrow E\left(S'^2\right) = E\left(S^2\right)^{\beta} \tag{12}$$

where β is estimated as the average ratio between S^1 and S^2 for the genuine scores.

Then, the final score is computed as the weighted sum of the two partial scores:

$$S = \alpha \cdot \beta \cdot S^2 + (10 - \alpha) \cdot S^1 \tag{13}$$

where $\alpha \in [0, 10]$ and β is the aforementioned normalising parameter.

This way, it follows that the final encrypted score $E(S)$ can be directly computed from the partial fingerprint and signature encrypted scores, $E(S^1)$ and $E(S^2)$, as follows:

$$E(S) = E\left(S^2\right)^{\alpha \cdot \beta} \cdot E\left(S^1\right)^{10-\alpha} \tag{14}$$

Seven steps are therefore carried out:

0. During enrolment, the reference biometric templates are encrypted using the server public key *pk*. The encrypted templates $E(\mathbf{T}_r^1)$ and $E(\mathbf{T}_r^2)$ (see Eqs. (8) and (11)) are stored in the database.
1. Biometric samples are acquired and two different templates, $\mathbf{T}_p^1$ and $\mathbf{T}_p^2$, are extracted on the client.
2. The server sends the encrypted enrolled templates to the client for each biometric characteristic, $E(\mathbf{T}_r^1)$ and $E(\mathbf{T}_r^2)$.
3. The client computes in parallel the similarity score for each biometric characteristic in the encrypted domain: $E(S^1)$ and $E(S^2)$, according to Eqs. (7) and (10) (depending on the distance measure selected as encrypted similarity function).
4. The client fuses the individual scores according to Eq. (14).
5. The final encrypted score $E(S)$ is sent to the server.
6. The server decrypts the score using the secret key, *sk*.
7. Finally, the server compares the score to the verification threshold δ and outputs a genuine/impostor verification decision.

4.3.3. Decision level fusion (Fig. 5)

As in the score level fusion, in this case, each probe biometric sample acquired at the client is processed separately, generating two separate templates: $\mathbf{T}_p^1$ and $\mathbf{T}_p^2$. Both templates are independently compared on the client to $E(\mathbf{T}_r^1)$ and $E(\mathbf{T}_r^2)$, generating two partial scores $E(S^1)$ and $E(S^2)$, which are sent to the server. The final binary decision is computed by the server taking into account both partial decisions (D^1 and D^2), fused with the OR rule:

$$D^1 = \left(S^1 > \delta^1\right) \tag{15}$$

$$D^2 = \left(S^2 > \delta^1\right) \tag{16}$$

$$D_{\mathrm{OR}} = D^1 \quad \mathrm{OR} \quad D^2 \tag{17}$$

Although the OR rule has been considered in this study, as the proposed protection framework is general, any other logic rule could also be used (e.g., AND).

As in the previous case, seven steps are carried out:

0. During enrolment, the reference biometric templates are encrypted using the server public key *pk*. The encrypted templates $E(\mathbf{T}_r^1)$ and $E(\mathbf{T}_r^2)$ (see Eqs. (8) and (11)) are stored in the database.
1. Biometric samples are acquired and two different templates, $\mathbf{T}_p^1$ and $\mathbf{T}_p^2$, are extracted on the client.
2. The server sends the encrypted enrolled templates to the client for each biometric characteristic, $E(\mathbf{T}_r^1)$ and $E(\mathbf{T}_r^2)$.
3. The client computes in parallel the similarity score for each biometric characteristic in the encrypted domain: $E(S^1)$ and $E(S^2)$, according to Eqs. (7) and (10) (depending on the distance measure selected as encrypted similarity function).
4. The individual encrypted scores, $E(S^1)$ and $E(S^2)$, are sent to the server.
5. The server decrypts both scores using the secret key, *sk*.
6. Each score is compared to its corresponding threshold (δ^1 and δ^2) in order to generate the individual decision, D^1 and D^2.

7. Finally, the server fuses the individual D^1 and D^2 decisions (in this particular case following the OR rule) and outputs a genuine/impostor verification decision.

5. Experimental protocol

In order to prove the soundness of the proposed scheme, it is necessary to assess each requirement established within the ISO/IEC IS 24745 on biometric information protection [10], namely: *i*) verification performance preservation, *ii*) irreversibility and *iii*) unlinkability. Additionally, in order to prove the efficiency of the system, it is also required to *iv*) study the increase in the computational load due to the operations carried out in the encrypted domain. To achieve all four goals, an experimental and theoretical analysis will be carried out, involving three key steps:

- **Performance evaluation**: verification performance will be evaluated in Section 6 for a particular case study on on-line signature and fingerprint fusion over the publicly available BiosecurID Multimodal database [29]. We will compare the performance of the unimodal and multimodal systems, for the unprotected and the protected scenarios. More details on the database and particular systems used are included in the next subsections.
- **Irreversibility and unlinkability analysis**: both properties of the protected templates will be theoretically analysed in Section 7 in order to ensure the privacy of the subjects.
- **Complexity analysis**: finally, we will study the computational complexity at verification time in Section 8, in terms of the most costly operations and the storage requirements.

5.1. BiosecurID multimodal database

It should be noted that most MBTP schemes have been evaluated over chimeric databases (i.e., the different characteristics associated to one chimeric identity come from different individuals) instead of real multimodal databases (i.e., all the characteristics of a real identity are provided by the same subject). This way, possible correlations among different biometric characteristics belonging to the same subject, which reduce the accuracy of the systems, are not taken into account. Such a behaviour is shown in the only reference providing such comparison [24]. Furthermore, an evaluation of the corresponding biometric systems using unprotected data is presented in very few cases, thereby preventing the assessment of the performance degradation due to the protection mechanism. To avoid such limitations, a multimodal database is used in the present article, being recognition performance evaluated for both protected and unprotected domains.

Very few multimodal databases including on-line signature data are available. Among them, BiosecurID DB [29] is one of the most recently acquired, comprising fingerprint, signature, face, hand, iris and speech data belonging to 400 subjects. All samples were acquired in four time-spanned sessions at six different sites in an office-like uncontrolled environment simulating a realistic scenario.

For the on-line signature subset, four genuine signatures were captured in each session with the Wacom Intuos3 A4 Inking Pen Tablet, thus yielding $400 \times 4 \times 4 = 6400$ genuine signatures. The fingerprint subset comprises data of four fingers per subject, captured with a thermal and an optical sensor (Biometrika FX2000). For the present study, only the right index acquired with the optical sensor has been considered, therefore having $400 \times 4 \times 4 = 6400$ fingerprint samples.

In order to establish a fair comparison between unimodal and multimodal performance, we have designed a common protocol for all three scenarios (i.e., on-line signature, fingerprint and multibiometrics). The database is divided into a train set (first 50 subjects) and a test set (last 350 subjects). The score normalization parameter β and the score fusion parameter α (see Eq. (14)) are estimated over the train set and performance is evaluated over the test set. Regarding the test set, the first 300 subjects are enrolled and modelled with the four samples captured in the first session. The remaining 12 samples of those first 300 individuals are used for computing the genuine scores ($12 \times 300 = 3600$ genuine scores). Then, the first sample of the last 50 subjects are compared to each user model, leading to $50 \times 300 = 15,000$ impostor scores.

Table 2
Performance evaluation. EERs for the unimodal and multibiomtric systems for the unprotected and the protected domains.

	Euclidean		Cosine	
	Unprotected	Protected	Unprotected	Protected
Signature	4.62	4.62	5.05	5.05
Fingerprint	1.59	1.59	3.04	3.04
Feature	0.12	0.12	3.00	3.00
Score	0.74	0.74	1.25	1.25
Decision	1.19	1.19	1.71	1.71

5.2. Unprotected biometric verification systems

In the proposed fusion framework, alignment-free fixed-length templates are required. Among the state-of-the-art signature and fingerprint verification systems, we chose the on-line signature verification system based on global features proposed in [58] and the FingerCodes representation for fingerprints [45].

5.2.1. On-line signature verification

In this particular system, signatures are parametrized using the optimum set of 40 global features found in [59] from the total 100 features proposed in [58]. Such features include information such as the total duration of the signature, the number of pen-ups or the average speed. Similarity scores are computed according to the two similarity measures (i.e., Euclidean and Cosine) described in Section 4.2.

5.2.2. Fingerprint verification

In the FingerCode scheme presented in [45], a region of interest is located and divided into 80 sectors. These sectors are filtered with eight Gabor filters, and the final template comprises the standard deviations of the grey values comprised by each sector for each filter. From the original $80 \times 8 = 640$ features, a subset of the best performing 100 has been selected with the method proposed in [60]. As in the signature case, similarity scores are computed using the two distance measures described in Section 4.2, with no specific pre-alignment between samples. The alignment of fingerprints can be also done according to the direction of the highest curvature point, which is used as reference point.

6. Performance evaluation

According to the ISO/IEC 24745 international standard [10], biometric template protection schemes should preserve the verification performance of their unprotected counterparts. In this section, we evaluate the performance of the proposed fusion schemes for the particular case of on-line signature and fingerprints fusion, following the protocol described in Section 5. The Detection Error Trade-Off (DET) curves of the multimodal systems are depicted in Fig. 6 for the two distances considered, in the unprotected (solid) and protected (dashed) domains. A summary of all the Equal Error Rates (EER) is shown in Table 2, where the EERs for the unimodal systems are also included.

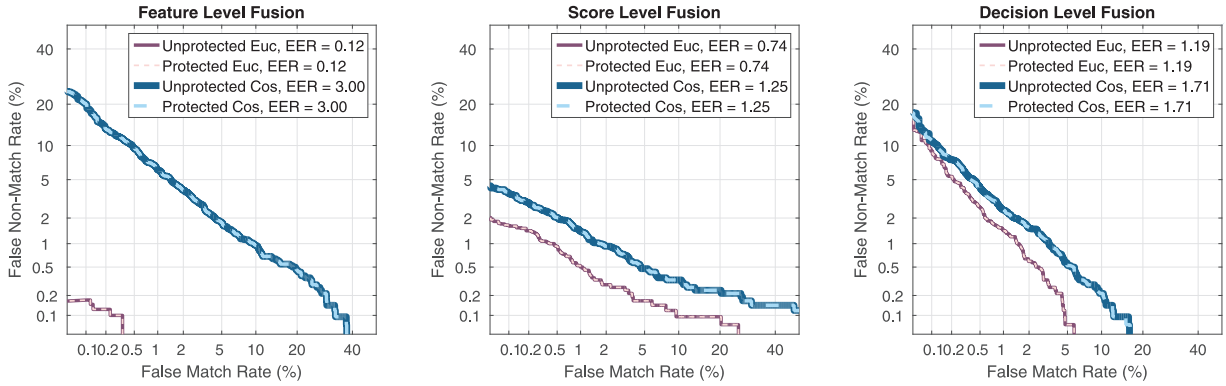

Fig. 6. Performance evaluation. DET curves for the Euclidean (thin purple) and the Cosine similarity (thick blue) for the unprotected (solid) and the protected (dashed) templates, for all the fusion approaches. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Regarding unimodal systems, both distances achieve the same performance in the plain and ciphertext domains for both characteristics. However, while they show almost the same performance for on-line signatures, for fingerprints the Euclidean distance performs better (EER of 1.59% vs 3.04%). The reason behind this difference lies in the feature selection step: the unprotected fingerprint system [45] was optimized to work with the Euclidean distance, thus yielding the best results with this metric.

For the multi-biometrics schemes, we first optimized the parameters α and β (see Eq. (14) for the score level fusion) over the train subset, using exhaustive search, in order to obtain the best possible performance. Then the performance was evaluated over the test subset. Two general trends may be observed in Fig. 6:

- The main take-away message of the performance evaluation is that there is no accuracy loss in the protected domain: at all fusion levels the Euclidean and Cosine distances are robust to the rounding errors introduced by HE.
- As a secondary observation, the Euclidean distance performs better in all fusion scenarios. This is a direct consequence of the performance for the unimodal systems, for which it shows a higher accuracy. In particular, for the Euclidean distance the EER decreases 92% at feature level, 53% at score level and 25% at decision level with respect to the best performing unimodal characteristic, the fingerprint (see Table 2).

7. Irreversibility and unlinkability analysis

As described in Section 1, in order to ensure the privacy of the subjects, biometric template protection systems should grant both irreversibility (i.e., no biometric information should be leaked by the protected template) and unlinkability (i.e., it should not be possible to cross-match protected templates from the same subject). We will analyse these two properties in the present section.

As mentioned in Section 3, for all multi-biometric fusion levels three different pieces of information should be hidden: *i*) only the client can have access to the plain probe biometric data $\mathbf{T}_p$, *ii*) the plain reference templates $\mathbf{T}_r$ should not be seen by any entity, being only their encrypted version $E(\mathbf{T}_r)$ stored, and *iii*) the plain score S should not be transmitted as it can potentially be used to perform hill-climbing or inverse-biometrics attacks.

For each distance measure considered, the information exchanged from the DB server to the client is the encrypted reference template $E(\mathbf{T}_r)$. Since only the authentication server knows the decryption key, *sk*, but he has never access to any protected or unprotected templates, there is no way for the client or any of the servers to learn any information from it. Conversely, the client sends no information about the acquired probe samples $\mathbf{T}_p$ to any server. Given that the decisional composite residuosity is an NP-hard problem, decoding the templates without *sk* could be considered computationally infeasible. We may thus conclude that the first requirement established by the ISO/IEC 24745 standard, irreversibility, is met.

Since irreversibility is granted, contrary to the case of unencrypted templates, no biometric information can be derived from stolen encrypted templates. However, they can still be used to impersonate the subject. In that case, a new key pair (*sk*, *pk*) could be generated and the entire database could be re-encrypted (i.e., re-secured) without having to re-acquire any new samples from the enrolled subjects, thereby achieving renewability (as pointed out in Section 2, this is not possible with cancelable biometrics approaches).

Finally, unlinkability is also granted. On the one hand, since unencrypted distances (i.e., similarity scores) between plaintexts are not preserved in the encrypted domain, given two samples $\mathbf{M}_1$ and $\mathbf{M}_2$ belonging to a given subject, their corresponding protected templates $E(\mathbf{T}_r^1)$ and $E(\mathbf{T}_r^2)$, encrypted with the same or different keys, are not related. On the other hand, since the Paillier cryptosystem provides semantic security against chosen-plaintext attacks [61], given a protected template $E(\mathbf{T}_r^1)$, no information can be feasibly derived about the original unprotected features $\mathbf{T}_r^1$. That way, no comparison can be established in the unprotected domain between some kind of information retrieved from the protected templates.

Furthermore, since the Paillier cryptosystem is based on probabilistic encryption, the randomness incorporated in the encryption algorithm leads to different ciphertexts given a particular message. This means that if $\mathbf{T}_r$ is encrypted twice with the same key, the corresponding ciphertexts could not be matched: $E_{pk_1}(\mathbf{T}_r, s_1) \neq E_{pk_1}(\mathbf{T}_r, s_2)$.

It should also be noted that, as stated in Section 4, only the server has access to the plain score S, and the only output is a genuine/impostor verification decision. Therefore, attacks based on the evolution of the score for different probe signatures, like the hill-climbing attacks described in [53,62], or the inverse biometrics methods proposed in [4,5], are prevented: they lack the necessary feedback to reconstruct an appropriate template or biometric sample.

Regarding each multi-biometric fusion level, templates are equally irreversible. However, the complexity level varies: since both score and decision levels require a separate storage of encrypted templates, feature level has been identified as the preferable approach [30,37]. Furthermore, while only one encrypted score is sent from the client to the server for the feature and score levels, decision level fusion in our approach requires the exchange of two different encrypted scores (one per characteristic), which

Table 3
Detailed complexity analysis. Number of encryptions / decryptions, and operations carried out during verification, as well as storage requirements, where F denotes the number of features of each modality used, N the number of modalities fused, $F_{fused} = F_1 + \cdots + F_N$, and M the number of samples used at enrollment.

		Euclidean distance	Cosine similarity
Unimodal	Enc/Dec	0/1	0/1
	Prod.	$3M \cdot F - 1$	$M \cdot F - 1$
	Exp.	$2M \cdot F$	$M \cdot F$
	Temp. size	$2M \cdot F + M$	$M \cdot F$
Feature	Enc/Dec	0/1	0/1
	Prod.	$3M \cdot F_{fused} - 1$	$M \cdot F_{fused} - 1$
	Exp.	$2M \cdot F_{fused}$	$M \cdot F_{fused}$
	Temp. size	$2M \cdot F_{fused} + M$	$M \cdot F_{fused}$
Score	Enc/Dec	0/1	0/1
	Prod.	$3M \cdot F_{fused} - 1$	$M \cdot F_{fused} - 1$
	Exp.	$2M \cdot F_{fused} + N$	$M \cdot F_{fused} + N$
	Temp. size	$2M \cdot F_{fused} + M$	$M \cdot F_{fused}$
Decision	Enc/Dec	0/N	0/N
	Prod.	$3M \cdot F_{fused} - N$	$M \cdot F_{fused} - N$
	Exp.	$2M \cdot F_{fused}$	$M \cdot F_{fused}$
	Temp. size	$2M \cdot F_{fused} + M$	$M \cdot F_{fused}$

increases slightly the complexity of the system as shown in the next section.

8. Complexity analysis

Finally, the computational cost is estimated in terms of the most complex operations carried out at verification time, namely: products and exponentiations. No encryptions or decryptions, which are the most costly operations, are carried out in the local client at verification time for any of the distances or fusion level approaches. On the server, only one (feature and score levels) or two (decision level) decryptions are needed to compute the final decision D. The encryption of the reference templates $E(\mathbf{T}_r)$ stored in the database is done during the enrolment, where we can assume that time or speed are not restricted. This way, fast verification is achieved.

We will first analyse the complexity of the unimodal systems, to develop later the analysis for the multibiometric scenarios. To that end, we should take into account three considerations:

- In order to verify an identity claim, we need to compute M single distances between the probe and each enrolled template:

$$E(S) = E\left(\sum_{j=1}^{M} S^j\right) = \prod_{j=1}^{M} E\left(S^j\right)$$

  where M is the number of enrolled templates. Therefore, the complexity of computing a single distance should be multiplied by M.
- In order to combine those individual scores we need to perform $M - 1$ additional products in the encrypted domain.
- Similarly, we need to store M templates for each subject.

Table 3 shows the detailed complexity analysis for each distance, for the unimodal and each multibiometric scenario. It should be noted that, for the estimation of the template size (and exchanged data), the size of the modulo $n = p \cdot q$ has to be taken into account: for a length of $|n|$ bits, ciphertexts will be $2|n|$ bits long. In order to achieve a security comparable to a state-of-the-art RSA, we have chosen a modulo of length $|n| = 1024$ bits [63]. In Table 3, the template size is measured in terms of the number of ciphertexts stored. Taking a key length of $|n| = 1,024$, each ciphertext comprises 2048 bits = 0.25 KB. It is thus enough to divide those figures by four in order to know the corresponding size in KB.

On the one hand, for the Euclidean distance, each score $E(S_{euc})$, see Eq. (7), involves $2F$ exponentiations (2 for each factor) and $3F - 1$ products (2 for each factor and $F - 1$ to combine all factors) for each of the M enrolled samples. With the additional products for the combination of the partial scores, the final number of operations is

$M(3F - 1) + (M - 1) = 3M \cdot F - 1$ products.

$2M \cdot F$ exponentiations.

Regarding the template size $E(\mathbf{T}_r)_{euc}$ (see Eq. (8)), the server has to keep in the database $M \cdot (2F + 1)$ ciphertexts.

One the other hand, for the encrypted cosine similarity $E(S_{\cos})$, defined in Eq. (10), the client computes F exponentiations (one for each factor) and $F - 1$ products to combine all factors. With the additional products for the combination of the partial scores, the final number of operations is

$M \cdot (F - 1) + (M - 1) = M \cdot F - 1$ products.

$M \cdot F$ exponentiations.

Regarding the template size $E(\mathbf{T}_r)_{\cos}$ (see Eq. (11)), the server has to keep in the database $M \cdot F$ ciphertexts.

Based on those computations, in the multibiometrics scenarios, where N characteristics are fused, we should take into account several observations. First of all, for all fusion scenarios, the template comprises now F_{fused} features instead of F, thus increasing its size accordingly (it depends linearly on F). The only difference between the feature level and the other two fusion levels is the storage as single template or as N separate templates, one for each characteristic.

Regarding the number of operations, for the feature level fusion, we will perform verification in the same way as in the unimodal case, but the templates handled will now comprise $F_{fused} = F_1 + \cdots + F_N$ features. Since all figures depend linearly on F, we just need to change F by F_{fused}.

At score level, we need to perform all the operations for each individual template. Then, $N - 1$ additional products and N exponentiations have to be carried out in order to fuse the scores yielded by each characteristic with their corresponding weights. Therefore, for the Euclidean distance the number of operations is

$$\begin{aligned}(3M \cdot F_1 - 1) + \cdots + (3M \cdot F_N - 1) + (N - 1) \\ = \left(3M \cdot F_{fused} - N\right) + (N - 1) \\ = 3M \cdot F_{fused} - 1 \text{ products.}\end{aligned}$$

$$\begin{aligned}(2M \cdot F_1) + \ldots + (2M \cdot F_N) + N \\ = 2M \cdot F_{fused} + N \text{ exponentiations.}\end{aligned}$$

On the other hand, for the cosine similarity we compute:

$$\begin{aligned}(M \cdot F_1 - 1) + \cdots + (M \cdot F_N - 1) + (N - 1) \\ = \left(M \cdot F_{fused} - N\right) + (N - 1) \\ = M \cdot F_{fused} - 1 \text{ products.}\end{aligned}$$

$$\begin{aligned}(M \cdot F_1) + \ldots + (M \cdot F_N) + N \\ = M \cdot F_{fused} + N \text{ exponentiations.}\end{aligned}$$

At decision level, we need to carry out all the operations for each individual template. Since both products and exponentiations depend linearly on F, we only need to substitute F by F_{fused}. Additionally, since N separate partial similarity scores $E(S_1), \ldots, E(S_N)$ are sent from the client to the server, N instead of one decryptions need to be performed in order to output the D verification decision.

It should be finally noted that in most biometric systems, the number of enrolled samples, M, or fused characteristic, N, are low, in most cases lower than ten. Therefore, since M, $N \ll F$, the number of products and exponentiations increases linearly with F_{fused},

Table 4
Complexity analysis for the feature level fusion, where $F_1 = 40$, $F_2 = 100$ and $M = 4$.

	Euc	Cos
Encryptions / Decryptions	0 / 1	0 / 1
Products	1679	559
Exponentiations	1120	560
$\mathbf{T}_p$ size	1.09 KB	
$E(\mathbf{T}_r)_{dist}$ size	200.25 KB	140 KB
Exchanged data	202.25 KB	142 KB
Time	202.25 KB	142 KB

for all fusion levels and distances, achieving a linear complexity of $\mathcal{O}(F_{fused})$.

Building upon the previous calculations, the complexity of the particular system evaluated in the present article is included in Table 4, where we considered $F_1 = 40$ and $F_2 = 100$ features and $M = 4$ enrolment samples. Since there is only a difference of one product or decryption between different fusion levels, only figures for the feature level fusion are shown.

On the other hand, since the same amount of encrypted information is being stored (in a single template for the feature level fusion, and in two different templates for the score and decision levels), the template size remains unchanged for all fusion levels. However, with respect to the unprotected templates, storage requirements are multiplied by 183 for the Euclidean distance and by 128 for the Cosine distance. This is due to the fact that we need to store more numbers in the encrypted system (see Eqs. (8) and (11)), and that each number needs 2048 instead of 16 bits. However, it should be noted that encrypted templates need between 140 KB and 200 KB, and can be therefore handled by most applications.

On the other hand, as the only computational difference between the fusion levels is the computation of a single score (feature level) or the computation of two separate scores (score and decision levels), which might be fused by the client (score level, see Eq. (14)) or by the server (decision level), the only difference is the computation of one more product on the client for the feature and score level fusions, and of one more decryption on the server for the decision level. Therefore, whenever it is possible to acquire all the samples at the same location, the feature level is preferred: it shows the best verification performance, and it is the most computationally efficient (only one template is stored and only one decryption is needed).

In terms of time, using Kun Liu's implementation of the Paillier cryptosystem in Java[2], and running the experiments in a machine with an Intel Core i7 with four 2.67 GHz cores, one comparison takes about 5×10^{-4} s. Similarly, in an identification task, R instead of a single comparison will be carried out for each probe sample, thus requiring $5 \times 10^{-4} \cdot R$ seconds. In particular, for a database comprising $R = 10^6$ records, a single identification would take approximately 8 min. Even if it should be noted that this is just an illustrative approximation (code should be optimized, separate servers for the DB and authentication need to be incorporated and GPUs can be used to accelerate the computations), we may conclude that the proposed scheme can be implemented in real time applications.

Finally, regarding the two distance measures considered, while the Euclidean distance presents a better performance (EER approximately 10% lower, as shown in Section 6), it requires twice as many exponentiations and thrice as many products as the cosine similarity. Additionally, the Cosine similarity requires a smaller template $E(\mathbf{T}_r)_{\cos}$ (40% smaller). However, in spite of such considerations, given the small time and storage requirements for verification purposes, the Euclidean distance would be preferred in most applications, in which recognition accuracy is of the utmost importance. On the other hand, for identification tasks over large databases, should time be a bigger constraint, the Cosine similarity would be preferred.

[2] http://www.csee.umbc.edu/~kunliu1/research/Paillier.html

9. Results summary

The main findings of the article can be summarised in the following:

- There is no performance loss in the protected domain. Furthermore, for the proposed scheme an EER as low as 0.12% is achieved for the feature level fusion, showing a 92% relative improvement with respect to the best performing individual characteristic. We may therefore conclude that the loss on accuracy due to the use of baseline systems with higher error rates than the current state-of-the-art is compensated fusing only two modalities.
- Only secure irreversible templates are stored in the server's database, hence achieving irreversibility.
- Templates are also unlinkable and renewability is achieved, thus fulfilling the requirements of the ISO/IEC IS 24745 [10].
- Since no plain information is shared, no biometric information is leaked, thereby preventing hill-climbing or inverse biometrics attacks [53,62].
- The proposed scheme can be deployed for real-time applications: no encryptions and only one decryption are performed on the server at verification time, templates require at most 200 KB and a comparison requires about 5×10^{-4} s.
- Feature level fusion is preferable to the other two levels, since it achieves a better performance and a unique template is generated for each subject. On the other hand, score level fusion is more flexible: it can be implemented in a distributed manner, where each client extracts one biometric sample and computes the corresponding similarity score. In that case, the score fusion would be carried out by the server, without each client having access to the other clients' scores.

10. Conclusions

We have proposed in this article the first general framework for multi-biometric template protection based on Homomorphic Encryption, where all the information, either stored in the database or exchanged between the client (issuing the identity claim) and the server (holding the database and verifying the identity claim), is encrypted. Different models have been described and analysed for the three fusion levels considered in the ISO/IEC TR 24722 on multimodal and other multi-biometric fusion [26], namely: feature, score and decision level.

Experiments were carried out on the on-line signature and fingerprint subcorpora of the publicly available BiosecurID multimodal database, following a clear protocol in order to make our research reproducible and allow future comparisons to other methods. The performance evaluation showed that verification can be carried out in the encrypted domain with no degradation. At the same time, the system fulfils the requirements established in the ISO/IEC 24745 IS on biometric information protection [10], as shown in the irreversibility and unlinkability analysis carried out. Given the low computational cost of the system (one decryption on the server side and no encryptions at verification time, and template sizes around 200 KB), and the good performance obtained (EER = 0.12%), we may conclude that the subject's privacy is protected in an efficient manner.

On the other hand, using Homomorphic Encryption for biometric template protection entails some limitations. For instance,

the implementation of more sophisticated schemes usually implies a higher computational load, or pre-aligned samples may be required. As a consequence, as future work lines, in order to make the protection scheme as general as possible, other matching functions and variable length templates will be considered and the code will be further optimized in order to reduce the time requirements, main drawback of the present scheme, specially for identification purposes. In addition, more complex score level fusions, such as those based on quality measures [64], will be studied. In that particular case, an implementation of SVMs within the Paillier cryptosystem needs to be developed and the signal quality measures could be combined in their plaintext form at the client, before sending the encrypted score to the server.

Acknowledgements

The research reported in this paper has been supported in part by the German Federal Ministry of Education and Research (BMBF) and by the Hessian Ministry of Science and the Arts within the Centre for Research in Security and Privacy (CRISP, www.crisp-da.de), project CogniMetrics (TEC2015-70627-R) from Spanish MINECO/FEDER and *CECABANK*. This work was conducted during a research stay at Biometrics and Multimedia Forensics Lab at Roma Tre University, Roma, Italy.

References

[1] A.K. Jain, Biometric recognition, Nature 449 (2007) 38–40.
[2] J. Galbally, R. Cappelli, A. Lumini, G.G. de Rivera, D. Maltoni, J. Fierrez, J. Ortega-Garcia, D. Maio, An evaluation of direct and indirect attacks using fake fingers generated from ISO templates, Pattern Recognit. Lett. 31 (2010) 725–732.
[3] R. Cappelli, D. Maio, A. Lumini, D. Maltoni, Fingerprint image reconstruction from standard templates, IEEE Trans. Pattern Anal. Mach. Intell. 29 (2007) 1489–1503.
[4] J. Galbally, A. Ross, M. Gomez-Barrero, et al., Iris image reconstruction from binary templates: an efficient probabilistic approach based on genetic algorithms, Comput. Vision Image Understanding 117 (10) (2013) 1512–1525.
[5] M. Gomez-Barrero, J. Galbally, A. Morales, et al., A novel hand reconstruction approach and its application to vulnerability assessment, Inf. Sci. 268 (2014) 103–121.
[6] A. Hadid, N. Evans, S. Marcel, J. Fierrez, Biometrics systems under spoofing attack: an evaluation methodology and lessons learned, IEEE Signal Process. Mag. 32 (5) (2015) 20–30.
[7] European Parliament, EU Regulation 2016/679 on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation)(2016). URL http://europa.eu/rapid/press-release_IP-15-6321_en.htm.
[8] C. Rathgeb, A. Uhl, A survey on biometric cryptosystems and cancelable biometrics, EURASIP J. Inf. Security 3 (2011) 1–25.
[9] V.M. Patel, N. Ratha, R. Chellappa, Cancelable biometrics: a review, IEEE Signal Process. Mag. 32 (5) (2015) 54–65.
[10] ISO/IEC JTC1 SC27 Security Techniques, ISO/IEC 24745:2011. Information Technology - Security Techniques - Biometric Information Protection, 2011.
[11] K. Simoens, B. Yang, X. Zhou, et al., Criteria towards metrics for benchmarking template protection algorithms, in: Proc. Int. Conf. on Biometrics, ICB, 2012, pp. 498–505.
[12] P. Tuyls, B. Skoric, T. Kevenaar (Eds.), Security with Noisy Data. On Private Biometrics, Secure Key Storage and Anti-Counterfeiting, Springer, 2007.
[13] P. Campisi (Ed.), Security and Privacy in Biometrics, Springer, 2013.
[14] T. Ignatenko, F. Willems, Biometric systems: privacy and secrecy aspects, IEEE Trans. Inf. Forensics Security 4 (4) (2009) 956–973.
[15] T. Ignatenko, F. Willems, Information leakage in fuzzy commitment schemes, IEEE Trans. Inf. Forensics Security 2 (5) (2010) 337–348.
[16] C. Fontaine, F. Galand, A survey of homomorphic encryption for nonspecialists, EURASIP J. Inf. Security 2007 (2007) 1–15.
[17] M. Barni, G. Droandi, R. Lazzeretti, Privacy protection in biometric-based recognition systems: a marriage between cryptography and signal processing, IEEE Signal Process. Mag. 32 (5) (2015) 66–76.
[18] C. Aguilar-Melchor, S. Fau, C. Fontaine, et al., Recent advances in homomorphic encryption: a possible future for signal processing in the encrypted domain, IEEE Signal Process. Mag. 30 (2) (2013) 108–117.
[19] J. Troncoso-Pastoriza, F. Perez-Gonzalez, Secure signal processing in the cloud: enabling technologies for privacy-preserving multimedia cloud processing, IEEE Signal Process. Mag. 30 (2) (2013) 29–41.
[20] S. Ye, Y.L.ad J. Zhao, S.S. Cheung, Anonymous biometric access control, EURASIP J. Inf. Security 2009 (2009) 1–17.
[21] M. Barni, T. Bianchi, D. Catalano, et al., A privacy-compliant fingerprint recognition system based on homomorphic encryption and fingercode templates, in: Proc. Int. Conf. on Biometrics: Theory Applications and Systems, BTAS, 2010, pp. 1–7.
[22] J. Bringer, H. Chabanne, A. Patey, Privacy-preserving biometric identification using secure multiparty computation: an overview and recent trends, IEEE Signal Process. Mag. 30 (1) (2013) 42–52.
[23] K. Nandakumar, A.K. Jain, Biometric template protection: bridging the performance gap between theory and practice, IEEE Signal Process. Mag. (2015) 1–12.
[24] A. Nagar, K. Nandakumar, A. Jain, Multibiometric cryptosystems based on feature-level fusion, IEEE Trans. Inf. Forensics Security 7 (1) (2012) 255–268.
[25] A. Ross, K. Nandakumar, A. Jain, Handbook of Multibiometrics, Springer, 2006.
[26] ISO/IEC JTC1 SC37 Biometrics, ISO/IEC TR 24722:2007. Information Technology – Multimodal and other multibiometric fusion, 2007.
[27] N. Poh, J. Kittler, A unified framework for biometric expert fusion incorporating quality measures, IEEE Trans. Pattern Analy. Mach. Intell. 34 (1) (2012) 3–18.
[28] C. Rathgeb, C. Busch, New Trends and Developments in Biometrics, InTech.
[29] J. Fierrez, J. Galbally, J. Ortega-Garcia, et al., BiosecurID: a multimodal biometric database, Pattern Anal. Appl. 13 (2009) 235–246.
[30] P.P. Paul, M. Gavrilova, Multimodal cancelable biometrics, in: Proc. Int. Conf. Cognitive Informatics Cognitive Computing, ICCI*CC, 2012, pp. 43–49.
[31] A.M.P. Canuto, F. Pintro, J.C. Xavier-Junior, Investigation fusion approaches in multi-biometric cancellable recognition, Expert Syst. Appl. 40 (2013) 1971–1980.
[32] K. Nandakumar, A.K. Jain, Multibiometric template security using fuzzy vault, in: Proc. Int. Conf. on Biometrics: Theory, Applications and Systems, BTAS, 2008, pp. 1–6.
[33] S. Cimato, M. Gamassi, V. Piuri, et al., Privacy-aware biometrics: design and implementation of a multimodal verification system, in: Proc. IEEE Ann. Conf. Computer Security Applications, CCSA, 2008.
[34] Y. Sutcu, Q. Li, N. Memon, Secure biometric templates from fingerprint-face features, in: Proc. Conf. on Computer Vision and Pattern Recognition, CVPR, 2007.
[35] C. Rathgeb, M. Gomez-Barrero, C. Busch, et al., Towards cancelable multi-biometrics based on adaptive bloom filters: a case study on feature level fusion of face and iris, in: Proc. Int. Workshop on Biometrics and Forensics, IWBF, 2015, pp. 1–6.
[36] A. Othman, A. Ross, On mixing fingerprints, IEEE Trans. Inf. Forensics Security 8 (1) (2013) 260–267.
[37] E. Kelkboom, X. Zhou, J. Breebaart, et al., Multi-algorithm fusion with template protection, in: Int. Conf. on Biometrics: Theory, Applications, and Systems, BTAS, 2009, pp. 1–8.
[38] C. Fang, Q. Li, E.C. Chang, Secure sketch for multiple secrets, in: Proc. Int. Conf. on Applied Cryptography and Network Security, 2010, pp. 367–383.
[39] A. Juels, M. Sudan, A fuzzy vault scheme, Designs Codes Cryptography 38 (2) (2006) 237–257.
[40] A. Juels, M. Wattenberg, A fuzzy commitment scheme, ACM Conf. Comput. Commun. Security (1999) 28–36.
[41] E. Verbitskiy, P. Tuyls, C. Obi, B. Schoenmakers, Key extraction from general nondiscrete signals, IEEE Trans. Inf. Forensics Security 5 (2) (2010) 269–279.
[42] R.L. Lagendijk, Z. Erkin., M. Barni, Encrypted signal processing for privacy protection: conveying the utility of homomorphic encryption and multiparty computation, IEEE Signal Process. Mag. 30 (1) (2013) 82–105.
[43] A.C.-C. Yao, How to generate and exchange secrets, in: Proc. Annual Symposium on Foundations of Computer Science, SFCS, 1986, pp. 162–167.
[44] P. Paillier, Public-key cryptosystems based on composite residuosity classes, in: Proc. EUROCRYPT, 1999, pp. 223–238.
[45] A.K. Jain, S. Prabhakar, L. Hong, S. Pankanti, FingerCode: a filterbank for fingerprint representation and matching, in: Proc. Conf. on Computer Vision and Pattern Recognition, CVPR, 1999.
[46] T. Bianchi, S. Turchi, A. Piva, et al., Implementing fingercode-based identity matching in the encrypted domain, in: Proc. Workshop on Biometric Measurements and Systems for Security and Medical Applications, BIOMS, 2010, pp. 15–21.
[47] Z. Erkin, M. Franz, J. Guajardo, et al., Privacy-preserving face recognition, in: Privacy Enhancing Technologies, Springer, 2009, pp. 235–253.
[48] A.-R. Sadeghi, T. Schneider, I. Wehrenberg, Efficient privacy-preserving face recognition, in: Int. Conf. on Information, Security and Cryptology, ICISC, Springer, 2010, pp. 229–244.
[49] M. Osadchy, B. Pinkas, A. Jarrous, B. Moskovich, SCiFI: a system for secure face identification, in: Proc. IEEE Symp. Security and Privacy, 2010, pp. 239–254.
[50] M. Blanton, P. Gasti, Secure and efficient protocols for iris and fingerprint identification, in: Proc. European Symposium on Research in Computer Security, ESORICS, 2011, pp. 190–209.
[51] K. Simoens, J. Bringer, H. Chabanne, S. Seys, A framework for analyzing template security and privacy in biometric authentication systems, IEEE Trans. Inf. Forensics Security 7 (2) (2012) 833–841.
[52] O. Goldreich, The Foundations of Cryptography – Vol 2, Cambridge University Press, UK, 2004.
[53] M. Gomez-Barrero, J. Galbally, J. Fierrez, Efficient software attack to multimodal biometric systems and its application to face and iris fusion, Pattern Recognit. Lett. 36 (2014) 243–253.
[54] S. Goldwasser, S. Micali, Probabilistic encryption, J. Comput. Syst. Sci. 28 (2) (1984) 270–299.

[55] T. ElGamal, A public key cryptosystem and a signature scheme based on discrete logarithms, IEEE Trans. Inf. Theory 31 (a) (1985) 469–472.
[56] M. Gomez-Barrero, J. Galbally, J. Fierrez, Implementation of fixed-length template protection based on homomorphic encryption with application to signature biometrics, in: Proc. Int. Workshop on Biometrics and Forensics (IWBF), 2016.
[57] A.K. Jain, K. Nandakumar, A. Ross, Score normalization in multimodal biometric systems, Pattern Recognit. 38 (2005) 2270–2285.
[58] M. Martinez-Diaz, J. Fierrez, R.P. Krish, J. Galbally, Mobile signature verification: feature robustness and performance comparison, IET Biom. 3 (2014) 267–277.
[59] J. Galbally, J. Fierrez, J. Ortega-Garcia, Performance and robustness: a trade-off in dynamic signature verification, in: Proc. Int. Conf. on Acoustics, Speech and Signal Processing, ICCASP, 2008.
[60] E. Maiorana, P. Campisi, A. Neri, Feature selection and binarization for on–line signature recognition, in: Proc. Int. Conf. on Biometrics, ICB, 2009, pp. 1219–1229.
[61] R. Anderson, Security Engineering: A Guide to Building Dependable Distributed Systems, 2001.
[62] E. Maiorana, G.E. Hine, P. Campisi, Hill-climbing attacks on multi-biometrics recognition systems, IEEE Trans. Inf. Forensics Security 10 (5) (2015) 900–915.
[63] D. Catalano, R. Gennaro, N. Howgrave-Graham, The bit security of paillier's encryption scheme and its applications, in: Proc. EUROCRYPT, 2001, pp. 229–243.
[64] J. Fierrez, D. Garcia-Romero, J. Ortega-Garcia, J. Gonzalez-Rodriguez, Adapted user-dependent multimodal biometric authentication exploiting general information, Pattern Recognit. Lett. 26 (16) (2005) 2628–2639.

Marta Gomez-Barrero received the Ph.D. degree in Electrical Engineering in 2016, from Universidad Autonoma de Madrid, Spain. She currently works as a postdoctoral researcher at the Hochschule Darmstadt, Germany. Her research interests are mainly focused on pattern recognition and security and privacy evaluation of biometric recognition systems.

Emanuele Maiorana received the Ph.D. degree in biomedical, electromagnetism and telecommunication engineering with European doctorate label from Roma Tre University, Italy, in 2009. He is currently a research engineer with the department of engineering of Roma Tre University. His research interests are in biometric recognition and digital image/signal processing.

Javier Galbally received the Ph.D. degree in Electrical Engineering in 2009, from Universidad Autonoma de Madrid, Spain. He currently works as a postdoctoral researcher at the Joint Research Centre from the European Commission. His research interests are mainly focused on pattern and biometric recognition related problems.

Patrizio Campisi (Ph.D.) is Full Professor at Roma Tre University. His research interests are secure multimedia communications and biometrics. He has been General Chair of IEEE WIFS 2015, and of the 12th ACM Workshop on Multimedia and Security 2010, and technical co-Chair of IEEE WIFS 2012.

Julian Fierrez received the MSc and the PhD degrees in telecommunications engineering from Universidad Politecnica de Madrid in 2001 and 2006, respectively. Since 2010 he is an Associate Professor at Universidad Autonoma de Madrid. His research interests include signal and image processing, pattern recognition, and biometrics.