IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 26 May 2025, accepted 12 July 2025, date of publication 21 July 2025, date of current version 25 July 2025.

Digital Object Identifier 10.1109/ACCESS.2025.3590989

RESEARCH ARTICLE

Cancelable Face Biometrics With Soft-Biometric Privacy Enhancement

PIETRO MELZI[1], HATEF OTROSHI SHAHREZA[2,3], CHRISTIAN RATHGEB[4], RUBEN TOLOSANA[1], RUBEN VERA-RODRIGUEZ[1], JULIAN FIERREZ[1], (Member, IEEE), SÉBASTIEN MARCEL[3,5], AND CHRISTOPH BUSCH[6]

[1]Universidad Autónoma de Madrid, 28049 Madrid, Spain
[2]École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland
[3]Idiap Research Institute, 1920 Martigny, Switzerland
[4]Hochschule Darmstadt, 64295 Darmstadt, Germany
[5]Université de Lausanne, 1015 Lausanne, Switzerland
[6]Norwegian University of Science and Technology, 7034 Gjøvik, Norway

Corresponding author: Christian Rathgeb (christian.rathgeb@h-da.de)

This work was supported in part by European Union's Horizon 2020 Research and Innovation Program under the Marie Skłodowska-Curie under Grant 860813—TReSPAsS-ETN; and in part by German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE. The work of Ruben Tolosana and Ruben Vera-Rodriguez was supported by INTERACTION under Grant PID2021-126521OB-I00 MICINN/FEDER. The work of Julian Fierrez was supported by BBforTAI under Grant PID2021-127641OB-I00 MICINN/FEDER. UAM members are also supported by Cátedra ENIA UAM-VERIDAS en IA Responsable (NextGenera- tionEU PRTR TSI-100927-2023-2).

ABSTRACT The storage of biometric data has raised significant privacy concerns, necessitating robust measures for secure storage. While traditional Privacy-Enhancing Technologies (PETs), like Cancelable Biometric (CB) schemes, excel at creating protected templates that fulfill criteria such as irreversibility and unlinkability, they often fail to preserve the privacy of soft-biometric information. To address this issue, we propose a hybrid technology that combines PETs, leveraging their different properties to comprehensively address multiple privacy requirements and enhance overall protection for biometric templates. In our approach, we integrate Multi Incremental Variable Elimination (Multi-IVE), a recent technology designed to remove soft-biometric information from biometric templates, with conventional CB schemes. We apply our hybrid technology to facial templates and assess the properties of the resulting protected templates. In the event of stolen secrets, the combination of Multi-IVE with CB schemes helps decrease the accuracy of estimating soft-biometric attributes without affecting recognition performance, compared to CB schemes alone.

INDEX TERMS Biometrics, cancelable biometrics, soft-biometrics, privacy enhancement, face recognition.

I. INTRODUCTION

Biometrics has become a promising field in classical image and video processing [1]. It finds applications in various sectors, including commerce, government, and forensics [2], [3], [4], [5]. The term *biometrics* refers to the automated recognition of individuals based on their unique biological and behavioural characteristics, such as face or finger topography, iris structure, and handwritten signature dynamics. These characteristics are used to create distinguishable and repeatable biometric features for recognition purposes [6]. Biometric recognition systems store collections of these features, known as *biometric templates*, which are derived from biometric samples provided by individuals. The collection and use of biometric data pose significant ethical concerns related to privacy, consent, and surveillance. Unlike passwords, biometric identifiers are immutable and deeply personal, making their misuse particularly invasive and difficult to remediate if compromised. Ethical challenges arise when individuals are unaware of or unable to opt

The associate editor coordinating the review of this manuscript and approving it for publication was Vincenzo Conti.

out of biometric data collection, especially in contexts like public surveillance or workplace monitoring. Furthermore, the potential for function creep, where data collected for one purpose is used for another, raises issues of trust and informed consent. Disparities in accuracy across demographic groups also risk reinforcing systemic biases, particularly in law enforcement and border control. As biometric technologies proliferate, robust ethical frameworks and regulatory safeguards are essential to prevent abuse and ensure that individual rights and freedoms are respected. Moreover, an attacker who compromises a biometric database may impersonate enrolled individuals to access the authentication system. Additionally, biometric templates can reveal sensitive personal information like health conditions, emotions, and soft-biometric attributes [7]. In some cases, biometric templates can even be used to reconstruct the original biometric samples provided by individuals [8], [9], [10], [11], [12], [13].

Numerous efforts have been made to address these privacy concerns and ultimately improve the use of biometric data. In 2016, the European Union introduced the General Data Protection Regulation (GDPR)[1] to safeguard individuals and regulate the processing of their personal data. The GDPR treats biometric data as sensitive information, establishing data protection principles to govern their use. Additionally, the ISO/IEC 24745 standard on biometric information protection [14], first introduced in 2011 and updated in 2022, provides guidelines for protecting biometric templates stored in biometric recognition systems, defining two fundamental privacy requirements:

- *Irreversibility:* it should be computationally difficult to reconstruct biometric samples similar to the original captured samples from the stored biometric templates. Irreversibility can be achieved by applying irreversible transformations or transformations that make use of secret parameters to biometric data. We highlight the importance of considering also the possibility of partial irreversibility of biometric templates. In fact, partial reconstructions of the original data may be sufficient to allow attackers to access the system, and reveal soft-biometric information of individuals.
- *Unlinkability:* it should be computationally difficult to determine if different biometric templates belong to the same individual or not. Unlinkability can be obtained by introducing some randomness with keys or random parameters in transformations that protect biometric data. When unlinkability is satisfied, compromised biometric templates can be revoked and substituted with new ones, providing template renewability.

While meeting these requirements and ensuring no adverse impact on the functionality of biometric recognition systems, it is important to maintain the performance level achieved with unprotected biometric data when processing protected biometric templates.

The ISO/IEC 24745 standard also outlines a general architecture for generating biometric templates that adhere to the privacy requirements of irreversibility and unlinkability. This architecture has given rise to several Privacy-Enhancing Technologies (PETs) known as Biometric Template Protection (BTP) schemes, which can be classified into two main categories: Cancelable Biometrics (CB) and biometric cryptosystems. CB schemes involve intentional and repeatable distortions of the original biometric data, achieved through specific transformations that enable comparisons within the transformed domain. On the other hand, biometric cryptosystems are designed to either bind or generate a digital key from biometric data [15]. For comprehensive surveys on this topic, the interested reader is referred to [16], [17]. It is important to note that CB schemes *a)* do not always guarantee the desired privacy requirements without significant degradation in recognition performance [15], and *b)* while protecting the overall information within biometric templates through complex transformations, cannot prevent the extraction of soft-biometric attributes when attackers gain access to secret transformation keys and biometric templates can be at least partially reverted [8].

Different categories of PETs have been proposed to address these issues, each one employing distinct approaches to enhance privacy [8]. To maintain recognition performance without deterioration (*a*), one approach is to apply Homomorphic Encryption (HE) to biometric templates. This cryptography-based PET enables comparisons in the encrypted domain that, once decrypted, yield results equivalent to those that would be obtained in the plaintext domain [18]. However, HE is computationally demanding and provides privacy assurances only when the algorithm's key remains confidential. To prevent the extraction of soft-biometric attributes from biometric templates (*b*), other PETs have been specifically designed to safeguard or remove soft-biometric information from biometric templates. This is crucial since extracting such information without user consent raises significant privacy concerns [19], and soft-biometric attributes are highly entangled within biometric templates. Consequently, researchers have introduced an additional privacy requirement for biometric templates:

- *Privacy of soft-biometrics:* the extraction of soft-biometric attributes from biometric templates for purposes different than the originally intended ones must be prevented.

While various PETs are designed to fulfill specific privacy requirements, none have achieved comprehensive coverage of all these requirements. The different categories of PETs are not mutually exclusive; on the contrary, they can be integrated, leveraging their distinct properties to collectively enhance the privacy of biometric templates [8]. Hybrid PETs hold promise for the future of biometric privacy enhancement, with some combination of BTP and HE schemes already proposed in the literature [20], [21]. To the best of our knowledge, no hybrid PET has focused specifically on the safeguarding of soft-biometric information within

[1] https://bit.ly/3y0kDkD

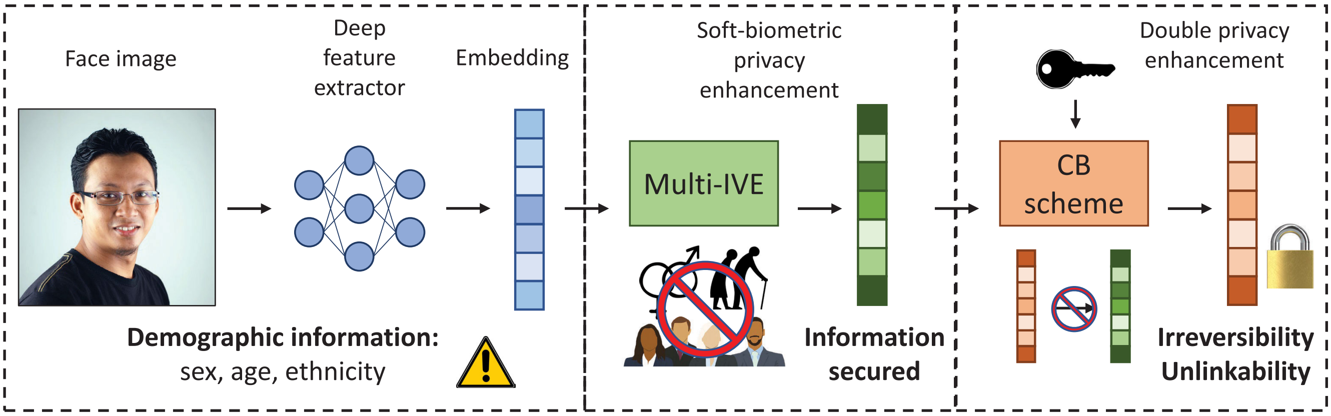

FIGURE 1. Both face images and their corresponding templates contain sensitive demographic information (*i.e,* soft-biometric attributes) requiring protection. While satisfying the requirements of irreversibility and unlinkability, Cancelable Biometric (CB) schemes are not designed to protect soft-biometric information within biometric templates. To address this issue, we propose a dual-layer protection strategy for biometric templates, combining soft-biometric privacy enhancement with CB schemes. Color image.

biometric templates. We hypothesize that combining a PET specifically designed for securing soft-biometric attributes with a CB scheme can provide several advantages. Therefore, in this study, we propose a hybrid PET obtained from the combination of Multi Incremental Variable Elimination (Multi-IVE), a novel PET based on the elimination of soft-biometric information from biometric templates [22], with three state-of-the-art CB schemes, namely *i)* BioHashing [23], *ii)* Multi-Layer Perceptron (MLP) Hashing [24], and *iii)* Index-of-Maximum (IoM) Hashing [25]. In Figure 1 we present an overview of our proposed combination of PETs. The obtained hybrid PET is expected to satisfy the following privacy requirements:

1) *Privacy of soft-biometrics:* the Multi-IVE component of our hybrid PET addresses this requirement by removing the information related to three demographic soft-biometric attributes, namely sex, age, and ethnicity. It is noteworthy that CB schemes alone cannot guarantee the removal of this information when their secret keys are exposed to attackers [8].
2) *Irreversibility and unlinkability:* CB schemes provide these properties to the biometric templates previously protected with Multi-IVE. The behavior of CB schemes should remain consistent regardless of the presence of Multi-IVE.
3) *Recognition performance:* when secret keys are unknown to attackers, the reduction in recognition performance caused by Multi-IVE can be offset by CB schemes employing user-specific and confidential keys. Our observations indicate that CB schemes can actually enhance recognition performance under these conditions [20], [26], [27], [28].

To comprehensively assess the properties of our proposed hybrid PET, we conduct evaluations in two distinct scenarios, each representing attackers with varying capabilities:

- *Normal scenario:* in this scenario, the secret keys of CB schemes are assumed to be unknown.
- *Stolen-key scenario:* in this scenario, it is assumed that attackers possess knowledge of the secret keys used in CB schemes.

In summary, our study introduces a hybrid PET that leverages the advantages of two PET categories: CB schemes and soft-biometric privacy enhancement. Our hybrid PET is expected to operate effectively in the normal scenario. However, if an attacker gains knowledge of secret keys, they may be able to partially revert biometric templates or reconstruct face image from protected templates [29]. In such scenarios, Multi-IVE ensures the security of demographic soft-biometric attributes[2] contained in the protected templates. We make the code available[3] for reproducibility of the experiments conducted in this study.

The remainder of the article is organized as follows. In Section II, we provide a description of the PETs that make up our hybrid PET, as well as other hybrid approaches aimed at enhancing the privacy of biometric templates. In Section III, we outline our proposed methodology, followed by the presentation of the experimental protocol employed in the study, which can be found in Section IV. We present and discuss the experimental results in Section V. Finally, we draw conclusions from our study in Section VI.

II. RELATED WORKS

A. SOFT-BIOMETRIC PRIVACY ENHANCEMENT

Soft–biometric PETs aim to address privacy concerns related to the use of biometric data by concealing selected soft–biometric attributes within it, such as sex, age, and ethnicity, and make it infeasible for the unsolicited

[2] We focus on the privacy-sensitive attributes that can be estimated given a biometric template, such as age, sex, ethnicity, from face templates. However, our approach is not limited to the mentioned attributes (age, sex, ethnicity) and can be used for any other attribute (e.g., skin tone, eye color, etc.) that can be estimated given a biometric template.

[3] https://github.com/PietroMelzi/HybridPET

extraction of sensitive personal information. Numerous soft-biometric PETs have been proposed in the literature. Given the increasing adoption of these techniques in real–world applications, it is essential to understand the extent to which attribute information can be recovered from privacy–enhanced biometric templates [30]. Soft–biometric PETs are designed with the explicit purpose of preventing the extraction of soft-biometric information from biometric data. This differs from CB schemes, which protect the entire biometric data without specific attention to the contained soft-biometrics attributes. Soft-biometric PETs typically employ one of the two approaches described in the following:

- *Information removal:* this approach involves identifying soft-biometric attributes within biometric data representations, removing them, and generating new representations of biometric data that exclude these soft-biometric attributes. An example of information removal is provided in [31].
- *Information protection:* according to this approach, the representation of biometric data is altered to prevent the extraction of soft-biometric attributes. Soft-biometric attributes are not discarded but are made inaccessible in the new representations of biometric data. Examples of information protection techniques are provided in [32], [33], [34].

1) MULTI-IVE

Multi-IVE is a soft-biometric PET designed to simultaneously safeguard multiple soft-biometric attributes within biometric templates derived from facial images. For technical details on Multi-IVE the interested reader is referred to [22]. This technology implements the previously described approach of information removal. It builds upon the original Incremental Variable Elimination (IVE) algorithm, which follows an iterative process of feature elimination from facial templates, ensuring a gradual reduction of the soft-biometric information contained therein [35]. The IVE algorithm is based on the training of a decision tree ensemble for the prediction of a selected soft-biometric attribute from facial templates. The trained decision tree ensemble is used to derive an importance measure for each feature in the template, quantifying the extent to which it provides information related to the selected soft-biometric attribute. The importance measure is used to identify and subsequently remove the features that carry the most substantial information about the selected attribute. The IVE algorithm is executed through multiple iterations, establishing the sequence for eliminating features within the facial templates to protect.

Multi-IVE employs Principal Component Analysis (PCA) to transform the facial template domain with the aim of simplifying information distribution and potentially gathering soft-biometric information into a reduced number of features. PCA is a linear transformation method that projects data into a space where orthogonal components capture the maximum variance in the data. Therefore, Multi-IVE focuses on the elimination of principal components within biometric templates, rather than individual features. Furthermore, it introduces innovative methods for the simultaneous detection and removal of multiple soft-biometric attributes from facial templates. Notably, Multi-IVE has demonstrated its effectiveness across various state-of-the-art feature extractors for facial images. The results from cross-database evaluations support its capacity to enhance the privacy of multiple soft-biometric attributes concurrently.

B. CANCELABLE BIOMETRIC (CB) SCHEMES

CB schemes are PETs utilized to create protected biometric templates that meet the privacy requirements of irreversibility and unlinkability. They achieve this by transforming an individual's original template, denoted as t^i for the individual i, into a protected template, denoted as $t_c^i = CB(t^i, k^i)$, where k^i corresponds to the user-specific keys. The protected template is distorted in such a manner that it becomes difficult to obtain the original biometric template [36]. The literature contains numerous studies that provide implementations of CB schemes and describe their properties [15], [37]. It is important to note that CB schemes are usually claimed to meet privacy requirements, but the extent to which these requirements are fulfilled can vary depending on the specific CB implementations and may be overestimated during evaluation. CB schemes also suffer from various security attacks as given in literature [36]. In this study, we employ three CB schemes that have been validated as meeting the requirements of both irreversibility and unlinkability, as demonstrated in the literature [20], [27]: namely *i)* BioHashing [23], *ii)* MLP Hashing [24], and *iii)* IoM Hashing [25].

1) BIOHASHING

BioHashing [23] is a CB scheme serving as a two-factor authentication method that combines tokenized random numbers with user-specific biometric templates. It operates by producing a user-specific orthogonal matrix and multiplying it to the unprotected templates. The outcome is subsequently binarized to generate binary-valued protected templates. In the recognition process, comparison scores are calculated using Hamming distance. BioHashing is robust against data capture offsets and provides security by requiring both the random data token and the user's biometric sample for code generation.

2) MULTI-LAYER PERCEPTRON HASHING

MLP Hashing [24] is a CB scheme that employs a MLP randomly initialized with the user-specific key. MLP Hashing processes unprotected templates through the MLP in a non-linear projection step, and subsequently binarizes the final MLP layer to produce the protected templates. Similarly to BioHashing, the recognition phase in MLP Hashing utilizes Hamming distance for comparing probe and reference templates. MLP Hashing has undergone evaluation based on unlinkability, irreversibility, and recognition accuracy to fulfill the ISO/IEC 30136 standard requirements. The results demonstrate its competitive performance, comparable to the BioHashing and IoM Hashing algorithms.

3) INDEX-OF-MAXIMUM HASHING

IoM Hashing [25] is a CB scheme that employs a ranking-based locality sensitive hashing to protect biometric templates. It applies a series of user-specific transformations to the unprotected templates and then derives IoM values for each transformation, which serve as the protected templates. Specifically, in this work we consider the variant known as Gaussian Random Projection-based (IoM-GRP), where user-specific Gaussian projection vectors are used for transformation. IoM-GRP Hashing generates integer-valued templates and employs the average number of collisions to calculate comparison scores during recognition. IoM Hashing provides robust concealment of biometric information, with strong assurance for irreversibility. Furthermore, it exhibits insensitivity to feature magnitude, making it resilient to variations in biometric features.

C. HYBRID PRIVACY-ENHANCING TECHNOLOGIES

The concept of hybrid PETs that combine the properties offered by different technologies has been explored in the literature, with hybrid PETs often combining different BTP schemes. Given the challenge of finding a single BTP method that can simultaneously provide both security and performance, researchers have proposed a hybrid approach that takes advantage of both biometric cryptosystems and CB schemes [38]. In this hybrid scheme, a three-step hybrid algorithm is proposed. The first step involves the creation of a cancelable template through a random projection (CB scheme). To maintain discriminability, a subsequent transformation is applied to compensate for the lost discriminating information, converting the cancelable template into a binary representation. Finally, a biometric cryptosystem scheme, *i.e.,* fuzzy commitment, is used to generate a protected template. The combination of CB schemes and biometric cryptosystems has proven successfully in several works [39], [40], [41]. Also, the combination of two biometric cryptosystems, such as fuzzy vault and fuzzy commitment scheme, has demonstrated improvements in both recognition performance and the security of fingerprint minutiae templates [42].

Novel instances of hybrid PETs combine Homomorphic Encryption with CB schemes [20], [21]. HE ensures the preservation of recognition performance, as it conducts operations in the encrypted domain and generates results that, once decrypted, match those of unencrypted operations. However, the security provided by HE schemes completely relies on the secrecy of the decryption key, and the computations performed in the encrypted domain tend to be computationally intensive. To solve these challenges, CB schemes can be applied to biometric templates prior to employing HE, to enhance both the security and privacy of the overall system. In fact, CB schemes provide irreversibility even in the event of HE key exposure and facilitate dimensionality reduction in biometric templates, enhancing computational efficiency within the encrypted domain [20]. A hybrid PET that combines the good properties of Bloom Filters (CB scheme) and HE is proposed in [21]. This technique ensures unlinkability and high recognition accuracy, while being about seven times faster than the traditional HE-based approach.

While offering notable advantages, these hybrid PETs do not incorporate any form of soft-biometric privacy enhancement when secret keys are compromised. To the best of our knowledge, the hybrid PET proposed in this study represents the first algorithm capable of enhancing soft-biometric privacy in biometric templates while also meeting the traditional requirements of irreversibility and unlinkability.

III. PROPOSED METHOD

We present our proposed hybrid PET using mathematical notation. In this study, we consider different sets containing biometric templates from distinct domains. These sets are denoted as T_U, T_C, T_M, and T_H, which respectively include unprotected templates t_u, templates t_c protected with CB schemes, templates t_m protected with Multi-IVE, and templates t_h protected with our hybrid PET. Our hybrid PET, denoted as $H : T_U \times K \times N \rightarrow T_H$, is applied to unprotected templates $t_u \in T_U$, which are extracted from facial images, to produce protected templates $t_h \in T_H$. Here, K denotes the set of user-specific keys employed by the CB scheme, while N represents the number of iterations performed by the Multi-IVE algorithm.

To quantitatively assess the presence of soft-biometric information sb within biometric templates t_x in a given domain X, we train Machine Learning (ML) classifiers C_X^{sb} to classify biometric templates t_x according to a specific soft-biometric attribute sb. We consider $sb \in \{s, a, e\}$, with $s =$ sex, $a =$ age, and $e =$ ethnicity, and $X \in \{U, C, M, H\}$. For each domain X and soft-biometric attribute sb of interest, we train a classifier C_X^{sb} using a development set of templates $t_x \in D_X$, and evaluate its performance with an evaluation set of templates $t_x \in E_X$, such that $C_X^{sb} : E_X \rightarrow [0, 1]$, $D_X \subset T_X$, $E_X \subset T_X$, and $D_X \cap E_X = \emptyset$. Finally, we consider a numerical threshold denoted as $max\_acc_{sb} \in [0, 1)$, representing the maximum acceptable accuracy when evaluating the classifier C_X^{sb} in the estimation of soft-biometric attribute sb.

When we train a classifier C_U^{sb} in the unprotected domain and assess its accuracy in predicting the soft-biometric attribute sb from unprotected templates $t_U \in E_U$, it is usually observed that $C_U^{sb}(E_U) > max\_acc_{sb}$. Traditionally, CB schemes $CB : T_U \times K \rightarrow T_C$ are applied to unprotected templates $t_u \in T_U$ to generate protected templates $t_c \in T_C$. This process can be defined as follows:

$$t_c = CB\left(t_u, k^i\right) \tag{1}$$

Here, $k^i \in K$ represents the user-specific key utilized in the CB scheme for individual i. In the *normal scenario*, k^i is assumed to be unknown. When we train a classifier C_C^{sb} in the domain obtained with the CB scheme and evaluate its accuracy in predicting the soft-biometric attribute sb from protected templates $t_c \in E_C$, we observe that $C_C^{sb}(E_C) < max\_acc_{sb}$, thanks to the secrecy of k^i. However, in the *stolen-key scenario*, it is assumed that the attacker possesses knowledge of k^i, making it ineffective as a safeguard for the soft-biometric attribute sb. Consequently, when we train a classifier C_C^{sb} and evaluate its accuracy using protected templates $t_c \in E_C$, we observe that $C_C^{sb}(E_C) > max\_acc_{sb}$.

For this reason, we employ the Multi-IVE algorithm, denoted as $MIVE : T_U \times N \to T_M$, on unprotected templates $t_u \in T_U$ to enhance the privacy of soft-biometric information and generate templates $t_m \in T_M$. We represent the first step of our hybrid PET as follows:

$$t_{m_n} = MIVE\left(t_u, n\right) \tag{2}$$

Here, $n \in N$ indicates the number of iterations performed by the Multi-IVE algorithm. For a fixed number of iterations n and within the domain obtained with Multi-IVE, we train a classifier C_M^{sb} and assess its accuracy in predicting the soft-biometric attribute sb from protected templates $t_{m_n} \in E_M$. If the number of iterations n is sufficiently high, we observe that $C_M^{sb}(E_M) < max\_acc_{sb}$.

At this point, instead of employing CB schemes on unprotected templates t_u as shown in (1), we can apply CB schemes to biometric templates t_{m_n} that have been already protected with Multi-IVE. This can be done also in the stolen-key scenario, as the estimation of the soft-biometric attribute sb remains below the predefined maximum acceptable threshold max_acc_{sb}, thanks to Multi-IVE. Additionally, we ensure the privacy guarantees offered by CB schemes when generating biometric templates with dual protection, denoted as t_h. We represent the second step of our hybrid PET as follows, with $H = CB \circ MIVE$:

$$t_h = CB\left(t_{m_n}, k^i\right) = CB\left(MIVE\left(t_u, n\right), k^i\right) = H\left(t_u, k^i, n\right) \tag{3}$$

We train a classifier C_H^{sb} in the double protected domain and evaluate its accuracy in predicting the soft-biometric attribute sb from protected templates $t_h \in E_H$. We observe that $C_H^{sb}(E_H) < max\_acc_{sb}$.

The two protection mechanisms CB and $MIVE$ are combined in a sequential manner. Multi-IVE is applied first, such that it only has to be trained once. Alternatively, an application of CB and $MIVE$ would require a more complex training and would decease the overall flexibility of the proposed system. To effectively utilize the proposed hybrid PET, it is imperative to ensure that the dual protection does not lead to a reduction in recognition performance for the protected templates t_h when compared to the performance obtained when working with unprotected templates t_u. We experimentally assess this aspect. In conclusion, it is important to highlight that the scenario with $n = 0$ corresponds to a situation where only CB schemes are employed on unprotected biometric templates t_u. In this specific case, we have $t_h = t_c$ and $t_m = t_u$. This observation is useful to evaluate the results obtained in the assessment of the proposed hybrid PET.

The proposed method is generic and could therefore be applied to any biometric data, given the availability of corresponding cancelable biometric and soft-biometric privacy-enhancement methods. The latter type of methods have solely been proposed for face images, which are also the main focus of this work. It is important to note that more sensitive information can be derived from faces compared to other biometric characteristics such as iris or fingerprints, see [43], [44]. Moreover, since the proposed method is applied at feature level, the computational complexity is low, which allows for a seamless integration into existing (face) biometric systems.

IV. EXPERIMENTAL SETTINGS

A. DATABASES

We consider the same public databases that were employed in both the training and evaluation phases of the Multi-IVE algorithm [22], to provide a better comparison of performance between our proposed hybrid PET and the original Multi-IVE algorithm. The training phase of Multi-IVE is crucial to determine the order of feature elimination in facial templates. For this purpose, we use the Color FERET database [45], which consists of 2,722 facial images obtained from 994 individuals, with demographic information about sex, age, and ethnicity. We categorize age into three intervals: 0-29, 30-49, 50+, and ethnicity into four categories: Asian, Black, White, and Others. We observe that approximately 63% of the images pertain to males, 47% to individuals aged 30 to 49 years, and 62% to individuals of white ethnicity. In the evaluation process we use the DiveFace database [32], which equally represents six classes obtained from the combination of sex (Male, Female) with three ethnic groups (Asian, Black, and White). To assess recognition performance and the accuracy of sex and ethnicity classification across various iterations of the algorithm, we choose a subset of 6,000 individuals equally representing the six demographic groups.

Since biometric recognition algorithms, in particular face recognition, are heavily based on ML and therefore prone to exhibit demographic bias [46]. Bias mitigation may be achieved through a fair composition of training data. Based on the determination of sources of bias across demographic groups, the optimal composition of such a training database may be investigated. To this end, synthetic image generation techniques have been found to be useful [47].

Moreover, the composition of the testing data on which the goodness of bias mitigation will be estimated turns out to be crucial. In this regard, focus could be put on the creation of a suitable testing database taking into account various factors which may cause demographic differentials in face recognition. These efforts should result in a training and testing platform designed to measure demographic bias and train/test mitigation approaches, which is beyond the scope of this work.

B. ALGORITHMIC DETAILS

In our study, we consider templates obtained with three state-of-the-art facial image feature extractors based on Deep Neural Networks, namely *i)* ArcFace [48], *ii)* MagFace [49], and *iii)* ElasticFace [50]. Each of these extractors requires facial images as input with dimensions of 112 × 112 pixels, generating facial templates composed of 512 features. To protect these templates, we employ our proposed hybrid PET, created by combining the Multi-IVE algorithm with CB schemes. Details of the implementation of the Multi-IVE algorithm, along with its possible variants, can be found in [22], which we recommend for a more detailed understanding of the algorithm. In our study, we implement Multi-IVE in the PCA domain, without imposing any constraint on the elimination of principal components. We opt for variant A of the algorithm, as we found that the other variants do not provide any significant performance improvements.

As for the CB schemes, we consider three state-of-the-art algorithms, namely *i)* BioHashing [23], *ii)* MLP Hashing [24], and *iii)* IoM Hashing [25]. For both the Multi-IVE algorithm and the three CB schemes, we maintain a consistent template size of 512 features. This ensures that we can make a fair performance comparison between Multi-IVE and our proposed hybrid PET. In the case of MLP Hashing, we consider a neural network with three layers of 1024, 1024, and 512 nodes, which enables the generation of templates with the desired dimensions. CB schemes require the utilization of user-specific keys for the protection of facial templates. In our study, we explore two key scenarios: the *normal scenario*, where a distinct key is assigned to each individual, modeling key secrecy, and the *stolen-key scenario*, where the same key is used for all individuals, assuming a key disclosure scenario.

C. EVALUATION

To evaluate the performance of our hybrid PET, we focus on two key metrics: accuracy of biometric recognition, and accuracy of sex and ethnicity attribute classification. These metrics are evaluated in both the *normal* and *stolen-key* scenarios, at various iterations n of the Multi-IVE algorithm. Initially, we focus on the facial templates extracted through the ArcFace method [48]. We provide dual protection to these templates with our hybrid PET, applying sequentially the Multi-IVE and BioHashing algorithms. To assess the validity of our hybrid PET, we provide a performance benchmark of the algorithm along two fundamental aspects of the algorithm:

- *Choice of CB scheme:* we compare BioHashing against the other CB schemes considered in our study while utilizing ArcFace for facial template extraction.
- *Choice of feature extractor:* we compare ArcFace against the other feature extractor considered in our study while securing facial templates with BioHashing as CB scheme.

To evaluate our hybrid PET, we replicate the experimental setup proposed in [22] for evaluating the Multi-IVE algorithm. We evaluate our hybrid PET at intervals of 30 iterations of the Multi-IVE algorithm, repeating the evaluation with 5 distinct seeds. The results are presented as the mean accuracy and standard deviation calculated from these five different executions, providing an insight into the evolving accuracy trend in recognition performance vs. estimation of soft-biometric attributes. Recognition performance is evaluated in the context of biometric verification, focusing on mated and non-mated comparisons of facial templates. The assessment of these comparisons is reported in terms of Equal Error Rate (EER). We employ a set of 6,000 individuals selected from the DiveFace database [32], each with three available images. Mated comparisons are derived from all feasible pairs of images belonging to the same individuals, while non-mated comparisons are obtained by matching an image of each individual with ten images from different random individuals.

To assess the accuracy in estimating soft-biometric attributes from protected facial templates $t_h \in E_H$, we employ ML classifiers tailored to the specific attributes of interest, and subsequently report their classification accuracy. In our case, the estimated attributes are sex and ethnicity, as both of them are represented in the DiveFace database. We continue to work with the same set of 6,000 individuals mentioned previously, equally representing these two soft-biometric attributes of interest. We select one image from each individual, and split the images into a development set D_H and an evaluation set E_H. These sets are divided with a 70% to 30% split ratio and a stratification based on sex and ethnicity to ensure balanced representation. As ML classifier C_H^{sb}, we opt for MLP, chosen for its consistent performance in the evaluation of Multi-IVE (variant A). MLP consistently exhibits the highest performance compared to other ML classifiers across different iterations of Multi-IVE. For further information on the used MLP-based classifier, the interested reader is referred to [22]. At intervals of 30 iterations, we train our MLP classifier, configured with a single hidden layer of 1,000 nodes. The learning rate adapts dynamically, with an initial value of 0.001, and we incorporate early stopping criteria based on lack of improvement after ten epochs, determined through a validation score computed with

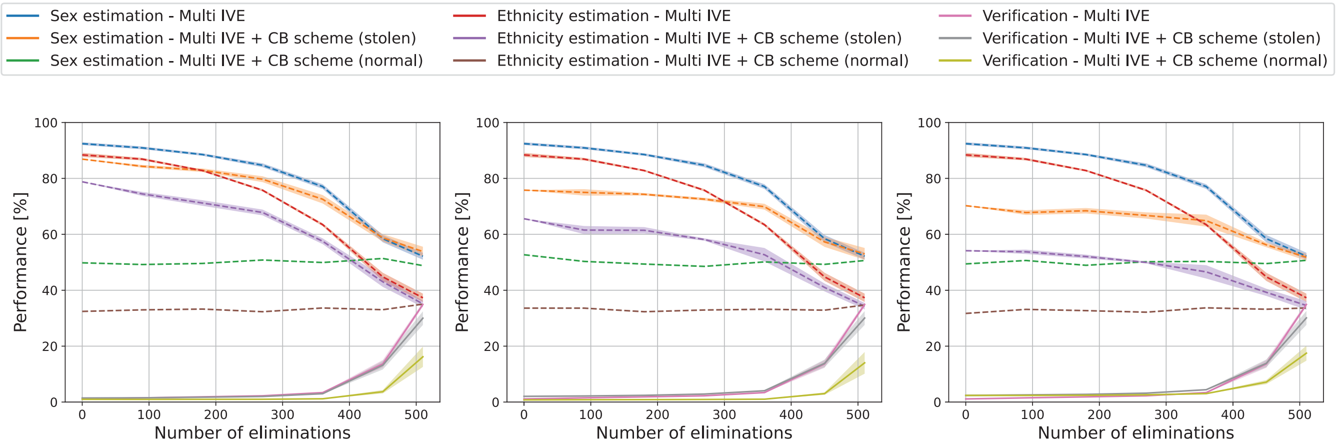

(a) CB scheme: BioHashing. Feature extractor: ArcFace

(b) CB scheme: MLP Hashing. Feature extractor: ArcFace

(c) CB scheme: IoM Hashing. Feature extractor: ArcFace

FIGURE 2. Evaluation of performance for different cancelable biometric (CB) schemes. We report performance in terms of: *i)* EER for recognition (solid lines), and *ii)* accuracy for soft-biometric estimation (dashed lines). Color image.

10% of the development set D_H. Subsequently, we evaluate the classifier's accuracy using the images allocated to the evaluation set E_H.

V. RESULTS

We present the results obtained during the evaluation of our proposed hybrid PET with a database different than the one used for training. In Figure 2 we benchmark the performance obtained for different CB schemes, while in Figure 3 we benchmark the performance obtained for different feature extractors. We evaluate the performance of our hybrid PET in terms of *i)* EER for biometric recognition, and *ii)* accuracy in estimating the sex and ethnicity attributes from protected templates $t_h \in E_H$.

Within each benchmark, we report every 30 iterations the performance provided by our hybrid PET with templates obtained considering different settings and scenarios: *i)* templates t_m originated solely from Multi-IVE, without the application of any CB scheme, *ii)* templates t_h generated with dual-layer protection, combining Multi-IVE and CB schemes in the *normal scenario*, and *iii)* templates t_h generated with dual-layer protection, combining Multi-IVE and CB schemes in the *stolen-key scenario*. Additionally, we closely observe the performance of CB schemes when Multi-IVE is not applied, *i.e.,* when the number of iterations performed by Multi-IVE n equals zero. This is illustrated in Figures 2 and 3 where the x-axis (labeled as "Number of eliminations") represents null values. In this specific context, the considered templates become respectively: *i)* unprotected templates t_u, *ii)* templates t_c protected with CB schemes in the *normal scenario*, and *iii)* templates t_c protected with CB schemes in the *stolen-key scenario*.

The accuracy of soft-biometric attribute estimation in the *normal scenario* approximates random values of 50% for sex and 33% for ethnicity, hence there is no need for averaging it across multiple seeds.

A. BENCHMARKING CB SCHEMES

In our first experiment, we evaluate the performance in recognition and soft-biometric attribute estimation provided by templates t_u, which have been originally extracted with ArcFace and subsequently protected with the Multi-IVE and BioHashing algorithms (Figure 2a). In the *normal scenario*, we obtained random accuracy when estimating sex and ethnicity attributes using our MLP classifier. Notably, the recognition performance exhibited by our hybrid PET in this scenario surpasses that of the Multi-IVE algorithm alone. As known from previous studies [20], [26], [27], the incorporation of secret user-specific keys in templates protected with BioHashing contributes to enhancing recognition performance.

With our proposed hybrid PET, our primary focus is on preventing the estimation of soft-biometric attributes from protected templates t_h in the *stolen-key scenario*. In this scenario, we observe that the dual protection offered by Multi-IVE and BioHashing results in a reduction of accuracy when estimating both sex and ethnicity, compared to the situations involving templates *i)* unprotected, *ii)* protected with Multi-IVE alone, and *iii)* protected with BioHashing alone. This is particularly evident until the elimination of the first 350 principal component of templates. Interestingly, we observe that the protection of soft-biometric attributes does not entail a simultaneous decrease in recognition performance. The latter is primarily governed by the number of iteration n performed by the Multi-IVE algorithm and results not affected by the combination of Multi-IVE with BioHashing. Therefore, within the *stolen-key scenario*, our hybrid PET provides an improvement in balancing the trade-off between recognition performance and soft-biometric protection, as the

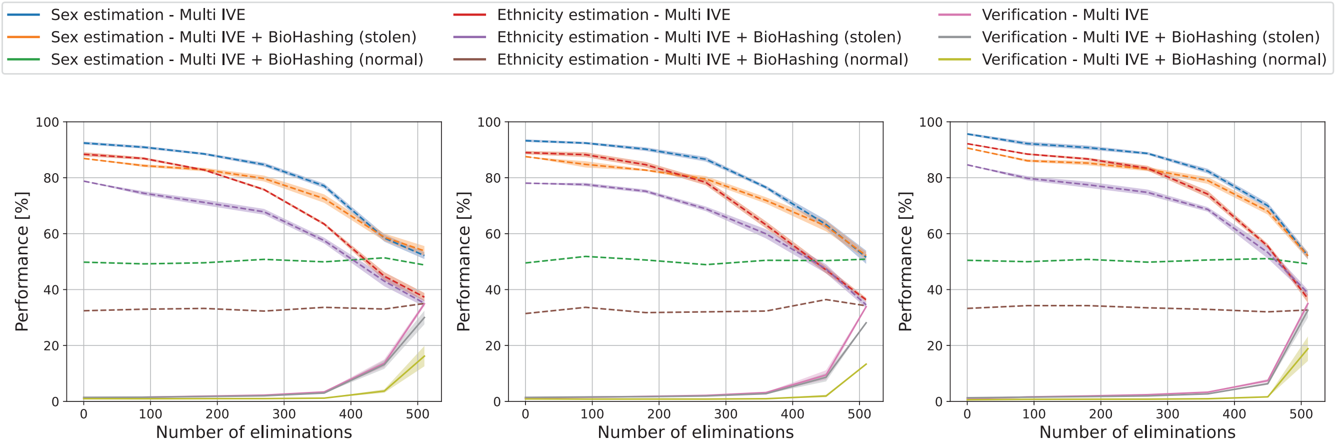

(a) Feature extractor: ArcFace. CB scheme: BioHashing.
(b) Feature extractor: MagFace. CB scheme: BioHashing.
(c) Feature extractor: ElasticFace. CB scheme: BioHashing.

FIGURE 3. Evaluation of performance for different feature extractors. We report performance in terms of: *i)* EER for recognition (solid lines), and *ii)* accuracy for soft-biometric estimation (dashed lines). Color image.

accuracy of sex and ethnicity estimation decreases without any concurrent increase of EER for recognition performance. This is a relevant outcome since the mentioned trade-off is usually considered a constraint for soft-biometric privacy enhancement [8].

Other CB schemes, *i.e.,* MLP Hashing and IoM Hashing, respect the trend observed for BioHashing in both the *normal* and *stolen-key* scenarios. Moreover, focusing on the *stolen-key* scenario when Multi-IVE is not applied (number of iterations $n = 0$), we observe that MLP and IoM Hashing provide lower accuracy in soft-biometric estimation compared to BioHashing. While templates protected with BioHashing achieve accuracy rates of 87% and 79% for sex and ethnicity estimations (Figure 2a), MLP Hashing provides accuracy rates of 76% and 66% (Figure 2b), and IoM Hashing provides accuracy rates of 70% and 54% (Figure 2c). Both MLP and IoM Hashing consistently maintain lower accuracy in soft-biometric estimation throughout the duration of the algorithm.

B. BENCHMARKING FEATURE EXTRACTORS

We also conduct benchmarks with three state-of-the-art feature extractors, as they represent soft-biometric information differently within the generated templates. However, the proposed hybrid PET provides comparable performance when applied to biometric templates extracted with different feature extractors. A subtle exception is observed in the case of ElasticFace (Figure 3c), where a slightly superior performance in soft-biometric estimation is noted. According to [51], the lower accuracy in ArcFace may be explained by the margin-principle used during training that distorts the feature space, making pattern learning harder. ElasticFace introduces an elastic margin loss that relaxes the fixed penalty margin of ArcFace and allows flexible space learning [50].

VI. CONCLUSION

In this study, we have introduced a novel hybrid PET designed to combine the strengths of soft-biometric privacy enhancement and CB schemes in protecting biometric templates extracted from facial images. The privacy concerns relative to the estimation of soft-biometric attributes from biometric templates are real, and they must be addressed during the protection of biometric templates. Our study represents the first integration in a hybrid approach of a technology intended to protect soft-biometric attributes within biometric templates with CB schemes. The results obtained during the evaluation confirm the validity of our hybrid approach, contributing to the improvement of the trade-off between recognition performance and protection of soft-biometric attributes.

In future works, HE schemes can be incorporated into the proposed hybrid PET to provide triplet protection. It has been shown that a hybrid method using CB and HE schemes is advisable. In fact, HE schemes provide theoretically proven security when the keys are not disclosed, while CB schemes apply a non-invertible transformation to the original templates and reduce the dimension of features prior to applying HE, which speeds up computations in the encrypted domain of HE [20]. Following these motivations and in combination with our proposed method, we can consider applying Multi-IVE prior to CB and HE schemes. In such a case, HE provides further security, and the soft-biometric attributes remain protected with Multi-IVE, even if the keys of HE and/or CB schemes are disclosed.

Our findings may be further generalized by considering alternative methods to evaluate the estimation of soft-biometric attributes. Finally, novel proposed attacks to the privacy of soft-biometric attributes need to be assessed during the evaluation of soft-biometric privacy enhancement [30], [52].

REFERENCES

[1] S. Bakshi, G. Guo, H. Proença, and M. Tistarelli, "IEEE access special section editorial: Visual surveillance and biometrics: Practices, challenges, and possibilities," *IEEE Access*, vol. 7, pp. 137638–137641, 2019.

[2] A. K. Jain, P. Flynn, and A. A. Ross, *Handbook of Biometrics*. Cham, Switzerland: Springer, 2007.

[3] J. Priesnitz, R. Huesmann, C. Rathgeb, N. Buchmann, and C. Busch, "Mobile contactless fingerprint recognition: Implementation, performance and usability aspects," *Sensors*, vol. 22, no. 3, p. 792, Jan. 2022.

[4] R. Tolosana et al., "SVC-onGoing: Signature verification competition," *Pattern Recognit.*, vol. 127, Jul. 2022, Art. no. 108609.

[5] I. Bouchrika, "A survey of using biometrics for smart visual surveillance: Gait recognition," in *Surveillance in Action*. Cham, Switzerland: Springer, 2018, pp. 3–23.

[6] *ISO/IEC 2382-37:2022 Information Technology—Vocabulary—Part 37: Biometrics*, Standard ISO/IEC JTC1 SC37, 2022.

[7] P. Delgado-Santos, G. Stragapede, R. Tolosana, R. Guest, F. Deravi, and R. Vera-Rodriguez, "A survey of privacy vulnerabilities of mobile device sensors," *ACM Comput. Surv.*, vol. 54, no. 11s, pp. 1–30, Jan. 2022.

[8] P. Melzi, C. Rathgeb, R. Tolosana, R. Vera-Rodriguez, and C. Busch, "An overview of privacy-enhancing technologies in biometric recognition," 2022, *arXiv:2206.10465*.

[9] H. O. Shahreza and S. Marcel, "Face reconstruction from facial templates by learning latent space of a generator network," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 36, 2023, pp. 12703–12720.

[10] H. O. Shahreza, V. K. Hahn, and S. Marcel, "Face reconstruction from deep facial embeddings using a convolutional neural network," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Oct. 2022, pp. 1211–1215.

[11] H. O. Shahreza and S. Marcel, "Comprehensive vulnerability evaluation of face recognition systems to template inversion attacks via 3D face reconstruction," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 12, pp. 14248–14265, Dec. 2023.

[12] H. O. Shahreza, V. K. Hahn, and S. Marcel, "Vulnerability of state-of-the-art face recognition models to template inversion attack," *IEEE Trans. Inf. Forensics Security*, vol. 19, pp. 4585–4600, 2024.

[13] H. O. Shahreza and S. Marcel, "Face reconstruction from partially leaked facial embeddings," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Apr. 2024, pp. 4930–4934.

[14] *Information Technology—Security Techniques—Biometric Information Protection*, Standard ISO/IEC 24745:2022, 2022.

[15] K. Nandakumar and A. K. Jain, "Biometric template protection: Bridging the performance gap between theory and practice," *IEEE Signal Process. Mag.*, vol. 32, no. 5, pp. 88–100, Sep. 2015.

[16] A. K. Jain, A. Ross, and U. Uludag, "Biometric template security: Challenges and solutions," in *Proc. 13th Eur. Signal Process. Conf.*, Sep. 2005, pp. 1–4.

[17] C. Rathgeb and A. Uhl, "A survey on biometric cryptosystems and cancelable biometrics," *EURASIP J. Inf. Secur.*, vol. 2011, no. 1, pp. 1–25, Dec. 2011.

[18] A. Acar, H. Aksu, A. S. Uluagac, and M. Conti, "A survey on homomorphic encryption schemes: Theory and implementation," *ACM Comput. Surv.*, vol. 51, no. 4, pp. 1–35, Jul. 2019.

[19] E. J. Kindt and E. J. Kindt, "Biometric data, data protection and the right to privacy," in *Privacy and Data Protection Issues of Biometric Applications: A Comparative Legal Analysis*. London, U.K.: Springer, 2013, pp. 87–272.

[20] H. O. Shahreza, C. Rathgeb, D. Osorio-Roig, V. K. Hahn, S. Marcel, and C. Busch, "Hybrid protection of biometric templates by combining homomorphic encryption and cancelable biometrics," in *Proc. IEEE Int. Joint Conf. Biometrics (IJCB)*, Oct. 2022, pp. 1–10.

[21] A. Bassit, F. Hahn, R. Veldhuis, and A. Peter, "Hybrid biometric template protection: Resolving the agony of choice between Bloom filters and homomorphic encryption," *IET Biometrics*, vol. 11, no. 5, pp. 430–444, Sep. 2022.

[22] P. Melzi, H. O. Shahreza, C. Rathgeb, R. Tolosana, R. Vera-Rodriguez, J. Fierrez, S. Marcel, and C. Busch, "Multi-IVE: Privacy enhancement of multiple soft-biometrics in face embeddings," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis. Workshops (WACVW)*, Jan. 2023, pp. 323–331.

[23] A. T. B. Jin, D. N. C. Ling, and A. Goh, "Biohashing: Two factor authentication featuring fingerprint data and tokenised random number," *Pattern Recognit.*, vol. 37, no. 11, pp. 2245–2255, Nov. 2004.

[24] H. O. Shahreza, V. K. Hahn, and S. Marcel, "MLP-hash: Protecting face templates via hashing of randomized multi-layer perceptron," in *Proc. 31st Eur. Signal Process. Conf. (EUSIPCO)*, Sep. 2023, pp. 605–609.

[25] Z. Jin, J. Y. Hwang, Y.-L. Lai, S. Kim, and A. B. J. Teoh, "Ranking-based locality sensitive hashing-enabled cancelable biometrics: Index-of-Max hashing," *IEEE Trans. Inf. Forensics Security*, vol. 13, no. 2, pp. 393–407, Feb. 2018.

[26] H. O. Shahreza, V. K. Hahn, and S. Marcel, "On the recognition performance of BioHashing on state-of-the-art face recognition models," in *Proc. IEEE Int. Workshop Inf. Forensics Secur. (WIFS)*, Dec. 2021, pp. 1–6.

[27] H. O. Shahreza, P. Melzi, D. Osorio-Roig, C. Rathgeb, C. Busch, S. Marcel, R. Tolosana, and R. Vera-Rodriguez, "Benchmarking of cancelable biometrics for deep templates," 2023, *arXiv:2302.13286*.

[28] H. O. Shahreza, Y. Y. Shkel, and S. Marcel, "Measuring linkability of protected biometric templates using maximal leakage," *IEEE Trans. Inf. Forensics Security*, vol. 18, pp. 2262–2275, 2023.

[29] H. O. Shahreza and S. Marcel, "Breaking template protection: Reconstruction of face images from protected facial templates," in *Proc. IEEE 18th Int. Conf. Autom. Face Gesture Recognit. (FG)*, May 2024, pp. 1–7.

[30] P. Rot, K. Grm, P. Peer, and V. Štruc, "PrivacyProber: Assessment and detection of soft–biometric privacy–enhancing techniques," *IEEE Trans. Dependable Secure Comput.*, vol. 21, no. 4, pp. 2869–2887, Jul. 2024.

[31] B. Bortolato, M. Ivanovska, P. Rot, J. Križaj, P. Terhörst, N. Damer, P. Peer, and V. Štruc, "Learning privacy-enhancing face representations through feature disentanglement," in *Proc. 15th IEEE Int. Conf. Autom. Face Gesture Recognit. (FG)*, Nov. 2020, pp. 495–502.

[32] A. Morales, J. Fierrez, R. Vera-Rodriguez, and R. Tolosana, "SensitiveNets: Learning agnostic representations with application to face images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 6, pp. 2158–2164, Jun. 2021.

[33] P. Terhörst, K. Riehl, N. Damer, P. Rot, B. Bortolato, F. Kirchbuchner, V. Struc, and A. Kuijper, "PE-MIU: A training-free privacy-enhancing face recognition approach based on minimum information units," *IEEE Access*, vol. 8, pp. 93635–93647, 2020.

[34] P. Delgado-Santos, R. Tolosana, R. Guest, R. Vera-Rodriguez, F. Deravi, and A. Morales, "GaitPrivacyON: Privacy-preserving mobile gait biometrics using unsupervised learning," *Pattern Recognit. Lett.*, vol. 161, pp. 30–37, Sep. 2022.

[35] P. Terhörst, N. Damer, F. Kirchbuchner, and A. Kuijper, "Suppressing gender and age in face templates using incremental variable elimination," in *Proc. Int. Conf. Biometrics (ICB)*, Jun. 2019, pp. 1–8.

[36] Manisha and N. Kumar, "Cancelable biometrics: A comprehensive survey," *Artif. Intell. Rev.*, vol. 53, no. 5, pp. 3403–3446, Jun. 2020.

[37] V. M. Patel, N. K. Ratha, and R. Chellappa, "Cancelable biometrics: A review," *IEEE Signal Process. Mag.*, vol. 32, no. 5, pp. 54–65, Sep. 2015.

[38] Y. C. Feng, P. C. Yuen, and A. K. Jain, "A hybrid approach for generating secure and discriminating face template," *IEEE Trans. Inf. Forensics Security*, vol. 5, no. 1, pp. 103–117, Mar. 2010.

[39] S. Kanade, D. Camara, E. Krichen, D. Petrovska-Delacretaz, and B. Dorizzi, "Three factor scheme for biometric-based cryptographic key regeneration using iris," in *Proc. Biometrics Symp.*, Sep. 2008, pp. 59–64.

[40] L. Leng and J. Zhang, "Dual-key-binding cancelable palmprint cryptosystem for palmprint protection and information security," *J. Netw. Comput. Appl.*, vol. 34, no. 6, pp. 1979–1989, Nov. 2011.

[41] W. J. Wong, M. L. D. Wong, and A. B. J. Teoh, "A security- and privacy-driven hybrid biometric template protection technique," in *Proc. Int. Conf. Electron., Inf. Commun. (ICEIC)*, Jan. 2014, pp. 1–5.

[42] A. Nagar, K. Nandakumar, and A. K. Jain, "A hybrid biometric cryptosystem for securing fingerprint minutiae templates," *Pattern Recognit. Lett.*, vol. 31, no. 8, pp. 733–741, Jun. 2010.

[43] A. Dantcheva, P. Elia, and A. Ross, "What else does your biometric data reveal? A survey on soft biometrics," *IEEE Trans. Inf. Forensics Security*, vol. 11, no. 3, pp. 441–467, Mar. 2016.

[44] A. Ross, S. Banerjee, and A. Chowdhury, "Deducing health cues from biometric data," *Comput. Vis. Image Understand.*, vol. 221, Aug. 2022, Art. no. 103438.

[45] P. J. Phillips, H. Moon, S. A. Rizvi, and P. J. Rauss, "The FERET evaluation methodology for face-recognition algorithms," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 10, pp. 1090–1104, Oct. 2000.

[46] P. Drozdowski, C. Rathgeb, A. Dantcheva, N. Damer, and C. Busch, "Demographic bias in biometrics: A survey on an emerging challenge," *IEEE Trans. Technol. Soc.*, vol. 1, no. 2, pp. 89–103, Jun. 2020.

[47] P. Melzi, C. Rathgeb, R. Tolosana, R. Vera-Rodriguez, A. Morales, D. Lawatsch, F. Domin, and M. Schaubert, "Synthetic data for the mitigation of demographic biases in face recognition," in *Proc. IEEE Int. Joint Conf. Biometrics (IJCB)*, Sep. 2023, pp. 1–9.

[48] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 4685–4694.

[49] Q. Meng, S. Zhao, Z. Huang, and F. Zhou, "MagFace: A universal representation for face recognition and quality assessment," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 14220–14229.

[50] F. Boutros, N. Damer, F. Kirchbuchner, and A. Kuijper, "ElasticFace: Elastic margin loss for deep face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2022, pp. 1577–1586.

[51] P. Terhörst, D. Fährmann, N. Damer, F. Kirchbuchner, and A. Kuijper, "Beyond identity: What information is stored in biometric face templates?" in *Proc. IEEE Int. Joint Conf. Biometrics (IJCB)*, Sep. 2020, pp. 1–10.

[52] D. Osorio-Roig, C. Rathgeb, P. Drozdowski, P. Terhörst, V. Štruc, and C. Busch, "An attack on facial soft-biometric privacy enhancement," *IEEE Trans. Biometrics, Behav., Identity Sci.*, vol. 4, no. 2, pp. 263–275, Apr. 2022.

PIETRO MELZI received the B.Sc. degree in engineering of computer systems and the M.Sc. degree in computer science and engineering from the Politecnico di Milano, Italy, in 2017 and 2020, respectively. He has pursued the Ph.D. degree with the Biometrics and Data Pattern Analytics (BiDA) Laboratory, Universidad Autónoma de Madrid, Spain, with a Marie Curie Fellowship within the TReSPAsS-ETN EU Project. His research interests include human–computer interaction, biometrics, privacy-enhancing technologies, computer vision, and pattern recognition.

HATEF OTROSHI SHAHREZA received the B.Sc. degree (Hons.) in electrical engineering from the University of Kashan, Iran, in 2016, the M.Sc. degree in electrical engineering from the Sharif University of Technology, Iran, in 2018, and the Ph.D. degree from the École Polytechnique Fédérale de Lausanne (EPFL), Switzerland, in 2024. He was a Research Assistant with the Idiap Research Institute, Switzerland, where he received the H2020 Marie Skodowska-Curie Fellowship for the Ph.D. degree. During the Ph.D. degree, he also spent six months as a Visiting Scholar with the Biometrics and Internet Security Research Group, Hochschule Darmstadt, Germany. He is currently a Postdoctoral Researcher with the Biometrics Security and Privacy Group, Idiap Research Institute. His research interests include deep learning, machine learning, computer vision, and biometrics. He was a recipient of European Association for Biometrics (EAB) Research Award 2023.

CHRISTIAN RATHGEB is currently a member with the Faculty of Computer Science, Hochschule Darmstadt (HDA), Germany. He is a Principal Investigator of the National Research Center for Applied Cybersecurity (ATHENE). His research interests include pattern recognition, iris and face recognition, security aspects of biometric systems, secure process design, and privacy enhancing technologies for biometric systems. He was a winner of the EAB—European Biometrics Research Award 2012, Austrian Award of Excellence 2012, the Best Poster Paper Awards (IJCB'11, IJCB'14, and ICB'15), the Best Paper Award Bronze (ICB'18), and the Best Paper Award (WIFS'21). He is a member of European Association for Biometrics (EAB), the Program Chair of the International Conference of the Biometrics Special Interest Group (BIOSIG), and an Editorial Board Member of *IET Biometrics* (IET BMT). He served for various program committees and conferences (e.g., ICB, IJCB, BIOSIG, and IWBF) and journals as a Reviewer (e.g., IEEE Transactions on Information Forensics and Security, IEEE Transactions on Biometrics, Behavior, and Identity Science, and IET BMT).

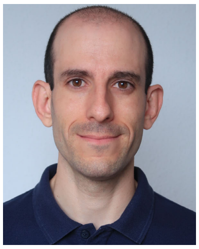

RUBEN TOLOSANA received the M.Sc. degree in telecommunication engineering and the Ph.D. degree in computer and telecommunication engineering from the Universidad Autónoma de Madrid, in 2014 and 2019, respectively. In 2014, he joined the Biometrics and Data Pattern Analytics (BiDA) Laboratory, Universidad Autónoma de Madrid, where he is currently an Assistant Professor. His research interests include signal and image processing, pattern recognition, and machine learning, particularly in the areas of DeepFakes, human–computer interaction, biometrics, and health. He is a member of the ELLIS Society, Technical Area Committee of EURASIP, and Editorial Board of the IEEE Biometrics Council Newsletter. He also received several awards, such as European Biometrics Industry Award (2018) from European Association for Biometrics (EAB) and the Best Ph.D. Thesis Award in 2019-2022 from Spanish Association for Pattern Recognition and Image Analysis (AERFAI). He served as the General Chair and the Program Chair (AVSS 2022) and the Area Chair (IJCB 2023 and ICPR 2022) in top conferences.

RUBEN VERA-RODRIGUEZ received the M.Sc. degree in telecommunications engineering from the Universidad de Sevilla, Spain, in 2006, and the Ph.D. degree in electrical and electronic engineering from Swansea University, U.K., in 2010. Since 2010, he has been with the Biometric Recognition Group, Universidad Autónoma de Madrid, Spain, where he has been an Associate Professor, since 2018. His research interests include signal and image processing, AI fundamentals and applications, HCI, forensics, and biometrics for security and human behavior analysis. He is actively involved in several national and European projects focused on these topics. He received the Medal in the Young Researcher Awards 2022 by Spanish Royal Academy of Engineering among other awards, and he has been a member of ELLIS Society, since 2023. He served as the Program Chair for some international conferences, such as: IEEE ICCST 2017, CIARP 2018, ICBEA 2019, and AVSS 2022. He has also organized several workshops and challenges in top conferences, such as WAMWB (MobileHCI 2023), KVC (BigData 2023), MobileB2C (IJCB 2022), and SVC-onGoing (ICDAR 2021).

JULIAN FIERREZ (Member, IEEE) received the M.Sc. and Ph.D. degrees from the Universidad Politécnica de Madrid, Spain, in 2001 and 2006, respectively. Since 2004, he has been with the Universidad Autónoma de Madrid, where he has been an Associate Professor, since 2010. His research interests include signal and image processing, AI fundamentals and applications, HCI, forensics, and biometrics for security and human behavior analysis. Since 2020, he has been a member of the ELLIS Society. He received best papers awards at AVBPA, ICB, IJCB, ICPR, ICPRS, and *Pattern Recognition Letters*; and several research distinctions, including EBF European Biometric Industry Award, in 2006, EURASIP Best Ph.D. Award, in 2012, the Miguel Catalan Award to the Best Researcher under 40 in the Community of Madrid in the General Area of Science and Technology, and the IAPR Young Biometrics Investigator Award, in 2017. He is also an Associate Editor of *Information Fusion*, IEEE Transactions on Information Forensics and Security, and IEEE Transactions on Image Processing.

SÉBASTIEN MARCEL received the Ph.D. degree in signal processing from CNET, research center of France Telecom (now Orange Labs), Université de Rennes I, France, in 2000. He heads the Biometrics Security and Privacy Group, Idiap Research Institute, Switzerland, and conducts research on face recognition, speaker recognition, vein recognition, attack detection (presentation attacks, morphing attacks, and deepfakes), and template protection. He is currently a Professor with the School of Criminal Justice, University of Lausanne, and a Lecturer with the École Polytechnique Fédérale de Lausanne. He is also the Director of Swiss Center for Biometrics Research and Testing, which conducts certifications of biometric products.

CHRISTOPH BUSCH is a member with Norwegian University of Science and Technology (NTNU), Norway. He holds a joint appointment with Hochschule Darmstadt (HDA), Germany. Further, he has been lecturing on biometric systems at DTU, Denmark, since 2007. On behalf of German BSI, he has been the coordinator for the project series BioIS, BioFace, BioFinger, BioKeyS Pilot-DB, KBEinweg, and NFIQ2.0. He was/is partner of the EU projects 3D-Face, FIDELITY, TURBINE, SOTAMD, RESPECT, TReSPsS, iMARS, and others. He is also the principal investigator in German National Research Center for Applied Cybersecurity (ATHENE) and is the Co-Founder of European Association for Biometrics (EAB). He is a member of the Editorial Board of the *IET Biometrics* and formerly of IEEE Transactions on Information Forensics and Security. Furthermore, he chairs the TeleTrusT biometrics working group as well as German standardization body on Biometrics and is convenor of WG3 in ISO/IEC JTC1 SC37.