

# Enhancing Privacy in Face Analytics Using Fully Homomorphic Encryption

Bharat Yalavarthi<sup>1,\*</sup> Arjun Ramesh Kaushik<sup>1,\*</sup> Arun Ross<sup>2</sup> Vishnu Boddeti<sup>2</sup> Nalini Ratha<sup>1</sup>

<sup>1</sup> University at Buffalo, New York, USA

<sup>2</sup> Michigan State University, Michigan, USA

{byalavar, kaushik3, nratha}@buffalo.edu, {rossarun, vishnu}@msu.edu

**Abstract**—Modern face recognition systems utilize deep neural networks to extract salient features from a face. These features denote embeddings in latent space and are often stored as templates in a face recognition system. These embeddings are susceptible to data leakage and, in some cases, can even be used to reconstruct the original face image. To prevent compromising identities, template protection schemes are commonly employed. However, these schemes may still not prevent the leakage of soft biometric information such as age, gender and race. To alleviate this issue, we propose a novel technique that combines Fully Homomorphic Encryption (FHE) with an existing template protection scheme known as PolyProtect. We show that the embeddings can be compressed and encrypted using FHE and transformed into a secure PolyProtect template using polynomial transformation, for additional protection. We demonstrate the efficacy of the proposed approach through extensive experiments on multiple datasets. Our proposed approach ensures irreversibility and unlinkability, effectively preventing the leakage of soft biometric attributes from face embeddings without compromising recognition accuracy.

## I. INTRODUCTION

Face recognition entails the extraction of features from face images and comparing them to either validate a claimed identity (“verification”) or determine an identity (“identification”) [1]. Recent advancements in deep neural networks and AI have resulted in the development of powerful face recognition systems [2], [3], [4] that can be deployed in a wide range of applications such as personalized services, law enforcement, border security, and smartphone access [5]. However, this development has also raised questions about privacy accorded to subjects and the security of the templates (such as embeddings) stored in a face recognition application [6]. Even as ethical concerns attendant to the technology are being rightfully discussed in public forums, it is necessary for the technology itself, on the one hand, and its users, on the other hand, to embrace measures that can enhance privacy and security while mitigating potential biases [7]. Otherwise, the technology runs the risk of being overwhelmed by restrictive legislation [8], [9] that can stifle the benefits of this technology in solving egregious crimes [10].

To address some of the privacy challenges associated with *face embeddings* stored as *templates*, we propose an approach in this paper that employs a polynomial transformation on homomorphically encrypted face embeddings.

\* First two authors contributed equally to this work.

Using Fully Homomorphic Encryption (FHE) in our proposed method ensures that the face recognition result is only disclosed to authorized parties with the secret key for homomorphic encryption, and the face embeddings themselves are secured through encryption during the recognition process. We show through our experiments that using FHE also prevents leakage of soft biometrics (e.g., age, gender/sex, race/ethnicity) from face embeddings. (It is necessary to point out that there is a difference between *race* and *ethnicity*, as well as *sex* and *gender*. In this paper, however, we use these terms interchangeably).

PolyProtect [11], a template protection scheme, transforms face embeddings into more secure templates using multivariate polynomials with user-specific parameters. Our research demonstrates a symbiotic relationship between FHE and PolyProtect in ensuring optimal security measures. The synergy between these two techniques is essential, as each contributes unique strengths that, when combined, establish a robust security framework. According to the inversion attack analysis conducted by PolyProtect [11], the risk of reversibility is notably high (30 - 99%) when an attacker possesses multiple (more than 4) templates of the same face. Even with the compromise of a single template, there is a 95% likelihood of reversibility when an *overlap* of 4 or greater is employed. However, PolyProtect introduces a tradeoff between accuracy and security, wherein increasing the *overlap* parameter enhances accuracy in face recognition and analysis tasks but concurrently diminishes template security. To address this challenge, we implement a strategy of encrypting the face embeddings using FHE and then applying the PolyProtect template with high overlap. This proactive step ensures that even under the full disclosure threat model, irreversibility, unlinkability and prevention of soft biometrics leakage are ensured without compromising identification performance.

Conversely, the integration of PolyProtect into our approach serves as a countermeasure against threats targeting compromised FHE systems, such as secret key leaks and passive attacks outlined in [12]. The identified risks pertain to FHE systems engaged in machine learning computations, such as mean and variance calculations. By combining both FHE and PolyProtect techniques (Fig. 1), our approach offers a more comprehensive and resilient privacy solution compared to relying on either method in isolation.

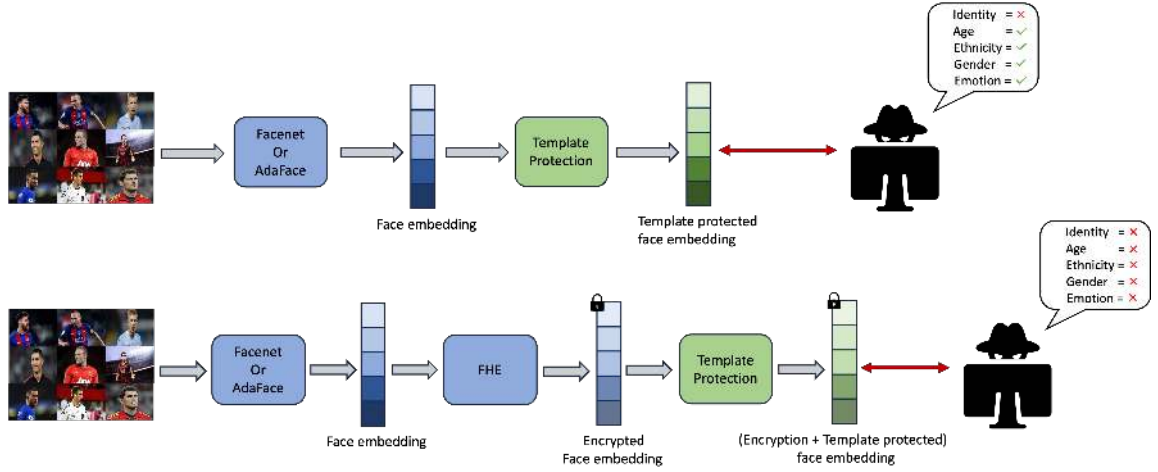


Fig. 1: An overview of our contribution in protecting face embeddings from leaking soft biometric attributes - Age, Gender, Ethnicity (FHE - Fully Homomorphic Encryption).

## II. THREAT MODEL

We use the term *facial analytics* to refer to the process of deducing semantic information from a face image. This could include sensitive information such as age, gender, ethnicity, and health [13] – sometimes known as soft biometrics. The possibility of deducing soft biometric cues from face images or their embeddings using automated techniques is a source of concern. These automated techniques can be machine learning models such as SVMs or deep neural networks (DNNs). For example, a face image or its embedding can be “stolen” by hackers and various soft biometrics can be derived from them thereby revealing sensitive information.

In our work, we presume that the face embedding is provided in an encrypted form. Our goal is to ensure that the encrypted embedding does not reveal any soft biometric information to unauthorized users. Note that the threat remains unchanged even if the parameters of the models used for extracting soft biometric information (e.g., weights of a DNN) are encrypted.

We consider the most challenging threat model according to ISO/IEC 30316, which is the full disclosure model, where the attacker possesses complete knowledge of the PolyProtect method [11], including its algorithm, number of embedding elements ( $m$ ), user-specific parameters (e.g.,  $C$ ,  $overlap$ , and  $E$ ), and one or more PolyProtected templates corresponding to a face embedding. In addition, we assume that the public key used in FHE is available but not the private key. If the embeddings are *not* encrypted, the hacker can infer soft biometric information from the PolyProtect template as we show in Table III.

## III. FHE BASICS

Encryption is the process where plain-text data is encrypted into ciphertext using a secret key and a cryptographic algorithm. Only authorized entities with a private key can decrypt the ciphertext back to the plaintext. Encryption is essential for protecting sensitive data from unauthorized

access or modification. Homomorphic Encryption (HE) is a cryptographic system that permits certain computations to be performed on encrypted data without requiring decryption [14]. In this system, we have public ( $pk$ ) and secret ( $sk$ ) keys, encryption ( $E$ ) and decryption ( $D$ ) mechanisms, and plaintext values  $x$  and  $y$ . When  $x$  and  $y$  are encrypted as  $x' = E(x, pk)$  and  $y' = E(y, pk)$ , respectively, a cryptosystem is considered homomorphic with respect to a chosen operator (e.g., addition or multiplication), denoted as  $\circ$ , if we can find another operator  $\bullet$  such that  $x \circ y = D(x' \bullet y', sk)$ . This means that we can conduct operations on encrypted data and obtain the same result when decrypting using the private secret key.

Specifically, given  $c_i = E(x_i, pk), i = 1, 2, \dots, K$ , an FHE scheme allows the computation of  $c = g(c_1, c_2, \dots, c_K)$  such that  $D(c, sk) = f(x_1, x_2, \dots, x_K)$  for any arbitrary function  $f$ .

It is essential to note that there are three types of homomorphic encryption schemes [15]: (a) Partial Homomorphic Encryption (PHE) permits addition or multiplication operations. (b) Somewhat Homomorphic Encryption (SHE) allows limited computations on ciphertexts. (c) Fully Homomorphic Encryption (FHE) enables computations on ciphertexts of any depth and complexity.

Numerous FHE systems have been introduced, including the BFV, BGV, and CKKS schemes [16]. The BFV and BGV schemes enable vector operations involving integers, while the CKKS scheme facilitates floating-point operations. These schemes achieve Single Instruction Multiple Data (SIMD) operations by bundling plaintext values into an array and then encrypting them to get ciphertext. In this work we use the HEAAN [17] library based on the CKKS scheme for FHE computations.

## IV. RELATED WORK

A notable body of literature explores privacy enhancements to soft biometrics at both the image and embedding (template) levels in face recognition. PFRNet [18] uses an Autoencoder framework to disentangle identity from attribute

information to suppress gender information in face embeddings. Similarly, SensitiveNets [19] uses a privacy-preserving neural network that suppresses soft biometrics attributes. The approach adopts an adversarial regularizer, which incorporates a sensitive information removal function into the learning objective. The Multi Incremental Variable Elimination (Multi-IVE) method [20] works by eliminating those feature variables in embeddings that predict soft biometric attributes. Increasing the number of eliminations was shown to decrease the soft biometrics leakage but significantly affected the identification performance. In [21], the authors introduce an adversarial attack approach designed to protect gender information in facial images. The method involves perturbing the image to minimize the estimated mutual information between the feature distribution acquired from a face recognition network and the gender variable. This method reduces gender leakage (prediction accuracy) by an average of 91% to 80% across three datasets. Reversible Attribute Privacy Preservation (RAPP) [22] uses a stream cipher to determine the sensitive attributes that have to be concealed with a user-defined password; it supports recovering the original attributes. It also uses an attribute adversarial network to generate perturbed images that conceal various attributes while retaining the utility of face verification. However, the identification performance is negatively impacted. The authors work around this challenge by reducing the number of features being concealed and the intensity of concealment. PrivacyNet [23] is another technique to impart soft biometric privacy to face images while preserving recognition capabilities via image perturbation using a GAN-based Semi-Adversarial Network (SAN). PrivacyNet also allows a person to choose the specific facial attributes to be obfuscated while allowing the other attributes to be extracted. One of the drawbacks of this image perturbation technique is that it sometimes does not generate realistic images and cannot conceal soft biometric features from a human observer.

Although current approaches mitigate the leakage of soft biometric attributes, they do not suppress it to the level of a “random guess”. In our work, we show that through the use of FHE, we can restrict this leakage in face embeddings to a level that is equal to or lower than that of a random guess. In [12], the authors show the susceptibility of privacy enhancement techniques such as homomorphic encryption. We address this susceptibility by employing a template protection scheme in addition to homomorphic encryption. Further, FHE encryption offers a stricter theoretical guarantee than existing methods for the security of soft biometrics.

FHE has been used in prior work for securing face recognition. Boddeti [24] proposes encrypting the face embeddings and performing face matching in the FHE domain. Batching and dimensionality reduction techniques are also explored to balance face-matching accuracy and computational complexity. In [25], the authors introduce an efficient approach for searching encrypted probe images against a large gallery, using fixed-length representations. In [26], the authors proposed a time-efficient and space-efficient face matching in the FHE domain for securing face templates.

Our approach stands out from previous works by integrating a template protection scheme, a compression technique, and FHE to enhance the security of face templates. Moreover, we conducted thorough experiments to assess their efficacy in mitigating the leakage of soft biometrics.

## V. TEMPLATE PROTECTION

Amongst the many existing protection templates, we have adopted **PolyProtect** [11] to showcase the benefits of our work. Let  $V = [v_1, v_2, \dots, v_n]$  denote an  $n$ -dimensional real-number face embedding. PolyProtect maps  $V$  to another real-number feature vector,  $P = [p_1, p_2, \dots, p_k]$  (where  $k < n$ ).  $P$  is the PolyProtected template of  $V$ .  $m$  (where  $m < n$ ) consecutive elements from  $V$  are mapped to single elements in  $P$  via a polynomial equation of coefficients,  $C = [c_1, c_2, \dots, c_m]$  and exponents,  $E = [e_1, e_2, \dots, e_m]$ .  $C$  and  $E$  are user-defined, non-zero, and distinct for each user of the face recognition system. The first  $m$  elements of  $V$  are mapped to  $p_1$  as :

$$p_1 = c_1 v_1^{e_1} + c_2 v_2^{e_2} + \dots + c_m v_m^{e_m} \quad (1)$$

Another important user-defined parameter is *overlap*, which defines the number of common elements from  $V$  used in successive values in  $P$ . When *overlap* = 0, the elements of  $V$  in each set are unique. The minimum and maximum values for *overlap* are 0 and  $m - 1$ , respectively. The mapping for  $p_2$  for overlaps 0 and  $m - 1$ , respectively, are as follows:

$$p_2 = c_1 v_{m+1}^{e_1} + c_2 v_{m+2}^{e_2} + \dots + c_m v_{m+m}^{e_m} \quad (2)$$

$$p_2 = c_1 v_2^{e_1} + c_2 v_3^{e_2} + \dots + c_m v_{m+1}^{e_m} \quad (3)$$

The authors of PolyProtect [11] have also performed an extensive survey on a number of existing Biometric Template Protection (BTP) methodologies and evaluated them based on recognition accuracy, irreversibility, and unlinkability [27]. According to the survey, none of the existing BTP methodologies before PolyProtect [11] satisfy all three criteria - recognition accuracy, irreversibility, and unlinkability.

## VI. ABLATION STUDY

### A. Variation of user-defined parameters ( $C$ , $m$ , *overlap*)

We conducted experiments with various user parameters in PolyProtect to investigate their impact on the leakage of soft biometrics. As *overlap* increased from 0 to 3, there was a notable increase in age leakage, particularly for values *overlap* > 0. However, gender and ethnicity showed consistent levels of leakage across different values of *overlap* (Fig. 2(a)). As the overlap increases, the amount of embedding information retained after the PolyProtect transformation also increases, potentially resulting in a higher risk of soft biometric leakage. In PolyProtect, the parameter  $m$  dictates the number of terms in the polynomial. At  $m = 6$ , we observed maximum leakage for age, whereas for gender and ethnicity, maximum leakage occurred at  $m = 7$ . Conversely, selecting  $m = 5$  minimized leakage across all three soft biometric attributes (Fig. 2(b)). Additionally, the parameter

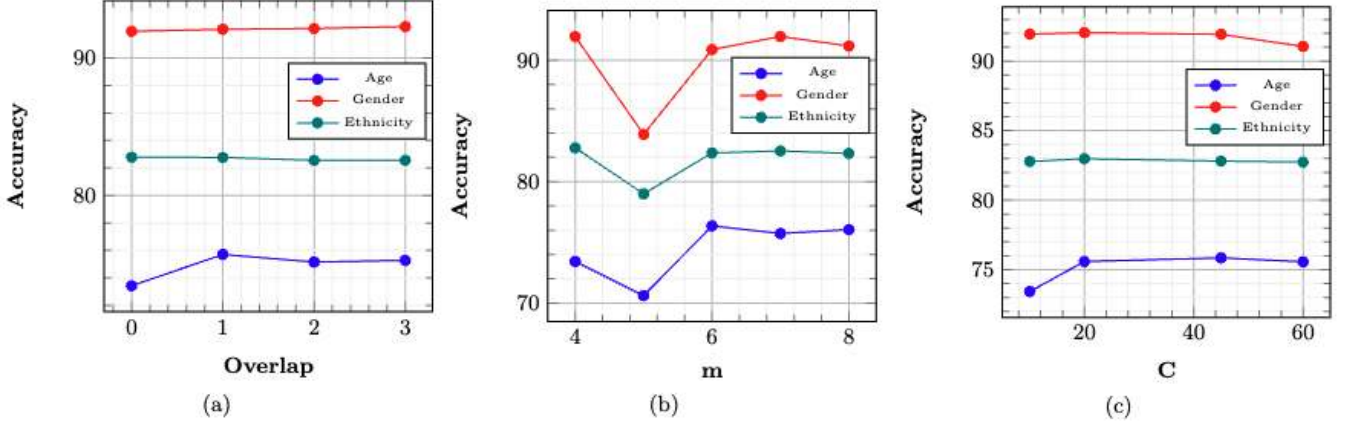


Fig. 2: Ablation study with the PolyProtect parameters shows soft biometrics leakage in different settings. (a) Overlap. (b)  $m$  - length of polynomial coefficients/exponents. (c)  $[-C, C]$  - range of polynomial coefficients.

$C$ , which defines the range of values for the polynomial coefficients  $[-C, C]$ , was varied from 10 to 60 in our study. The leakage of ethnicity remained relatively stable across all tested values of  $C$ , while age exhibited an increase from  $C = 10$  to 20, and gender showed a slight decrease from  $C = 45$  to 60 (Fig. 2(c)).

### B. Summation of Ciphertext elements

Implementing PolyProtect, cosine similarity, and fully connected layers in FHE requires summing up the elements within a ciphertext. This is not straightforward as we cannot access individual elements of a ciphertext. The three approaches to efficiently achieve the summation are described below:

1) **Naive Rotation**: This is a brute-force method for summation within a ciphertext where expensive ciphertext rotations are performed  $N - 1$  times and the running sum is computed until all the elements are covered (Algorithm 1).

2) **Discrete Fourier Transform**: When the Discrete Fourier Transform (DFT) of a signal is computed, the first value of DFT or the DC component will give the sum of the input signal values. We use this property to calculate the DFT of the ciphertext and get the sum of its values.

3) **Fold and Add**: This is a more efficient version of the naive rotation method described earlier, which can be visualized as iteratively folding the array into half and adding the corresponding folded parts  $\log_2 N - 1$  times as described in algorithm 2. Fig. 3 shows that the *Fold and Add* method was the fastest among the three with a significant speedup than the naive rotation method, especially for large ciphertext sizes.

### C. Polynomial approximation of inverse square root

Since computing cosine similarity requires inverse square root function, we use a polynomial approximation of the inverse square root function in the encrypted domain owing to the limitations of FHE in implementing non-linear functions. We restrict the input to the range  $(0, 1]$  to achieve a closer approximation. The performance of this approximation

**Algorithm 1** Add Ciphertext Elements Through Left Rotation  $n-1$  times

```

function NAIVE_ADD(Ciphertext  $c$ , long  $N$ )
  Ciphertext  $c1 \leftarrow c$ 
  while  $N > 0$  do
     $LeftRot(c1, 1)$   $\triangleright$  Left Rotates Ciphertext by 1
     $c \leftarrow Add(c, c1)$   $\triangleright$  Adds two Ciphertexts
     $N \leftarrow N - 1$ 
  end while
  return  $c$ 
end function

```

**Algorithm 2** Add Ciphertext Elements Through Left Rotation  $\log_2 N$  times

```

function FOLD_ADD(Ciphertext  $c$ , long  $N$ )
   $k \leftarrow \lceil \log_2 N \rceil$ 
   $i \leftarrow k - 1$ 
  while  $i > 0$  do
    Ciphertext  $c1 \leftarrow LeftRot(c, 2^i)$ 
     $c \leftarrow Add(c, c1)$ 
     $i \leftarrow i - 1$ 
  end while
  return  $c$ 
end function

```

is measured through the relative error of 2000 random points in the range  $(0, 1]$ . We consider both 6-degree and 8-degree polynomials as a tradeoff between computational depth and accuracy, as in Fig. 4.

## VII. PROPOSED APPROACH

The proposed methodology begins with the reception of encrypted face templates from users. These templates are generated through a pre-trained face recognition model and encrypted locally at the user's end using Fully Homomorphic Encryption (FHE), employing a private key exclusively known to the user. Subsequently, a template protection algo-

TABLE I: Statistics of the CelebSet Dataset (80 Identities).

Gender		Age				Ethnicity			
Males	Females	0-22	23-40	41-59	60+	Hispanic	White	Black	Asian
38,080	34,409	5,279	43,357	22,781	1,072	738	57,873	13,414	464
52.50%	47.50%	7.28%	59.82%	31.43%	1.47%	1.01%	79.85%	18.50%	0.64%




Fig. 3: Execution time to compute summation of elements within a ciphertext.




Fig. 4: Inverse square root function. (a) 6-degree polynomial approximation and (b) 8-degree polynomial approximation and their relative error over 2000 random points in the range (0,1].

rithm such as PolyProtect is applied to augment the security of the encrypted template, forming an additional layer of protection. The resulting encrypted PolyProtect template can be securely stored at the user's end and utilized as needed for tasks such as face identification or the extraction of soft biometrics, including gender, age, and ethnicity. The

outcomes of these tasks are transmitted to the user in encrypted form and can only be decrypted using the user's private key.

#### A. Dataset

We have performed our experiments on CelebSet [28] and Balanced Faces in the Wild (BFW) [29]. The statistics of these datasets have been detailed in Tables I, II. As the BFW dataset lacks age annotations for its face images, we utilized a pre-trained model trained on the CelebSet dataset to predict the ages of the BFW images. These predicted ages were then employed in our experiments.

#### B. Embedding Compression

It is commonly believed that compressing embeddings can improve privacy leakage. We explore embedding compression through **Matryoshka Representation Learning (MRL)** [30] as a technique to improve privacy, in addition to template protection and encryption. MRL is an innovative approach that enhances representation learning by encoding information at various granularities within a single embedding. MRL achieves this by capturing details at different levels of abstraction, enabling the model to adapt to downstream tasks efficiently. The method focuses on learning coarse-to-fine representations, ensuring that high-level, generalized features are initially captured and progressively refined with finer details. MRL stands out by delivering representations that are comparable in accuracy and rich in information when compared to independently trained low-dimensional representations, making it a promising solution for efficient and effective representation learning.

In our research, we leveraged MRL to compress face embeddings obtained from FaceNet/AdaFace, which are originally of 512-dimension, down to 64-dimension. As depicted in Fig. 5, we noted a decrease in classification accuracies when the compressed dimensions fell below 64. It is important to note that, since HEAAN supports SIMD operations, compression does not aid in reducing computational depth.

#### C. Face Identification

Given a vector of face embeddings of size  $n$ , and user-defined parameters  $C$ ,  $m$ , and  $overlap$ , we pad the embeddings vector such that it can be split into  $m$  small vectors  $V_i$  (where  $m \ll n$ ). Each  $m$ -sized vector is encrypted in FHE. We apply the PolyProtect algorithm as described in [11] on each  $V_i$ , resulting in vector  $P_i$ . Unlike in the PolyProtect algorithm [11], we store every PolyProtect mapping as a vector of  $m$  elements which are the same, resulting from the sum obtained through *Fold and Add* (Algorithm 2). We have employed the Cosine similarity as a metric to calculate



TABLE II: Statistics of the BFW Dataset (100 Identities).

Gender		Predicted Age						Ethnicity			
Males	Females	0-4	5-12	13-19	20-39	40-59	60+	Indian	White	Black	Asian
10,000	10,000	0	50	16,326	3,612	12	0	5,000	5,000	5,000	5,000
50%	50%	0%	0.25%	81.63%	18.06%	0.06%	0%	25%	25%	25%	25%




Fig. 5: Performance of Matryoksha Representation Learning (MRL) in extracting features - Identity, Age, Gender, Ethnicity - from different compressed dimensions. The graph is based on experiments performed using AdaFace on the CelebSet dataset.

similarity scores in the  $1 : N$  search for identification. Cosine similarity in FHE is computed using the algorithm described in algorithm 3.

**Algorithm 3** Cosine Similarity between ciphertexts  $C1$  and  $C2$  of size  $N$  in FHE

```

function COSINE_DISTANCE( $C1, C2, N$ )
   $C \leftarrow \text{Mult}(C1, C2)$ 
   $C \leftarrow \text{FOLD\_ADD}(C, N)$ 
   $D1 \leftarrow \text{FOLD\_ADD}(\text{Mult}(C1, C1), N)$ 
   $D2 \leftarrow \text{FOLD\_ADD}(\text{Mult}(C2, C2), N)$ 
   $D \leftarrow \text{Mult}(D1, D2)$ 
   $C \leftarrow \text{Normalization}(C)$   $\triangleright$  Scales all values to the
  range  $[0,1]$ 
   $D \leftarrow \text{Normalization}(D)$ 
   $D \leftarrow \text{approxInverseSquareRoot}(D)$ 
   $C \leftarrow \text{Mult}(C, D)$ 
  return  $C$ 
end function

```

#### D. Age, Gender and Ethnicity Prediction

We perform our experiments using various combinations of embeddings, template protection, compression, and encryption. The embeddings are extracted through FaceNet [31] and AdaFace [32] models. Our experiments are primarily aimed at determining the identification accuracy and soft biometrics leakage for the described combinations to prove

the efficacy of our proposed solution. Our networks (Fig. 6) have been tuned from the lens of an attacker to achieve maximum leakage of soft biometric features. For cases that require training with encrypted data in the FHE domain, we train an SVM classifier on the ASCII dump of the ciphertexts.

Apart from a myriad of leakage metrics defined in [33] and [34], we have chosen to define leakage through Suppression Rate (SR) and Privacy Gain (PG) [35]. Privacy Gain is defined as -

$$\text{PG} = (1 - R_p) - (1 - R_o) \quad (4)$$

where,  $R_o$  and  $R_p$  represent the recognition performances on the original data and the privacy-enhanced data, respectively. A positive value of PG signifies enhanced data protection. The ideal value for PG is 1.

Suppression Rate measures the difference attribute prediction accuracy with and without privacy enhancement,  $A_p$  and  $A_o$ , respectively

$$\text{SR} = (A_o - A_p) / A_o \quad (5)$$




Fig. 6: Feed Forward Network used for classifying soft biometrics.

## VIII. RESULTS

As evident from Tables III, IV, V our proposed approach prevents the leakage of soft biometrics from face embeddings with minimal loss in identification accuracy ( $< 2.5\%$ ), high Privacy Gain and high Suppression Rate. We can observe that our approach reduced the classification accuracies of soft biometric attributes to the level of random chance across the two datasets. We could also achieve an almost ideal Privacy Gain and Suppression Rate in certain scenarios, but we believe this could be because of the imbalanced nature of our datasets. Our experiments showcase that a combination of MRL, FHE, and PolyProtect in this order yields maximum protection against soft biometric leakage (Fig. 7 and Fig. 8).

The consistent efficacy of our method across two different embeddings (FaceNet and AdaFace) and datasets (CelebSet and BFW) shows the ability of our method to generalize in

TABLE III: Face identification and soft biometric classification accuracy (MRL - Matryoksha Representation Learning; FHE - Fully Homomorphic Encryption). Note that the proposed approach retains identification accuracy while successfully reducing soft biometric classification accuracy.

Embeddings	Dataset	Template Protection	Identification Accuracy	Gender Accuracy	Age Accuracy	Ethnicity Accuracy
FaceNet	CelebSet	None	99.42%	98.12%	87.68%	98.81%
		PolyProtect	99.42%	97.35%	85.00%	98.06%
		MRL	97.93%	98.00%	85.87%	97.38%
		MRL + PolyProtect	97.00%	96.32%	87.32%	98.44%
		MRL + FHE	97.83%	<b>52.22%</b>	<b>6.12%</b>	<b>8.01%</b>
		MRL + PolyProtect + FHE	96.95%	<b>52.22%</b>	<b>6.12%</b>	8.03%
	BFW	None	85.06%	95.25%	94.40%	91.97%
		PolyProtect	85.04%	95.07%	93.97%	91.62%
		MRL	84.42%	92.02%	93.98%	91.65%
		MRL + PolyProtect	84.31%	90.00%	87.90%	87.70%
		MRL + FHE	84.38%	49.98%	28.00%	24.01%
		MRL + PolyProtect + FHE	84.28%	<b>49.50%</b>	<b>27.82%</b>	<b>24.00%</b>
AdaFace	CelebSet	None	99.41%	96.93%	90.25%	96.31%
		PolyProtect	99.38%	95.75%	90.18%	96.56%
		MRL	97.95%	97.63%	85.50%	98.94%
		MRL + PolyProtect	97.59%	96.82%	89.13%	98.25%
		MRL + FHE	97.91%	<b>52.22%</b>	6.12%	8.02%
		MRL + PolyProtect + FHE	97.55%	<b>52.22%</b>	<b>6.11%</b>	<b>8.01%</b>
	BFW	None	88.55%	89.97%	85.02%	81.72%
		PolyProtect	88.46%	84.07%	84.62%	74.62%
		MRL	88.33%	87.55%	83.83%	76.43%
		MRL + PolyProtect	88.29%	85.58%	83.66%	75.58%
		MRL + FHE	88.23%	49.98%	27.97%	24.02%
		MRL + PolyProtect + FHE	88.21%	<b>49.50%</b>	<b>27.90%</b>	<b>24.00%</b>

TABLE IV: Privacy Gain across different soft biometric attributes.

Embeddings	Dataset	Template Protection	Identification Accuracy	Gender Attribute	Age Attribute	Ethnicity Attribute
FaceNet	CelebSet	PolyProtect	99.42%	0.72	2.68	0.75
		MRL	97.93%	0.12	1.81	1.43
		MRL + PolyProtect	97.00%	1.80	0.36	0.37
		MRL + FHE	97.83%	<b>45.90</b>	<b>81.56</b>	<b>90.80</b>
		MRL + PolyProtect + FHE	96.95%	<b>45.90</b>	<b>81.56</b>	90.78
	BFW	PolyProtect	85.04%	0.18	0.43	0.35
		MRL	84.42%	3.23	0.42	0.32
		MRL + PolyProtect	84.31%	5.25	6.50	4.27
		MRL + FHE	84.38%	45.27	66.40	67.96
		MRL + PolyProtect + FHE	84.28%	<b>45.75</b>	<b>66.58</b>	<b>67.97</b>
AdaFace	CelebSet	PolyProtect	99.38%	1.18	0.07	-0.25
		MRL	97.95%	-0.7	4.75	-2.63
		MRL + PolyProtect	97.59%	0.11	1.12	-1.94
		MRL + FHE	97.91%	<b>44.71</b>	84.13	88.29
		MRL + PolyProtect + FHE	97.55%	<b>44.71</b>	<b>84.14</b>	<b>88.30</b>
	BFW	PolyProtect	88.46%	5.90	0.40	7.10
		MRL	88.33%	2.42	1.22	5.29
		MRL + PolyProtect	88.29%	4.39	1.36	6.14
		MRL + FHE	88.23%	39.99	57.05	57.70
		MRL + PolyProtect + FHE	88.21%	<b>40.47</b>	<b>57.12</b>	<b>57.72</b>

different conditions. We also prove the ability of FHE to work seamlessly with template protection schemes and embedding compression techniques for preserving the privacy of face templates.

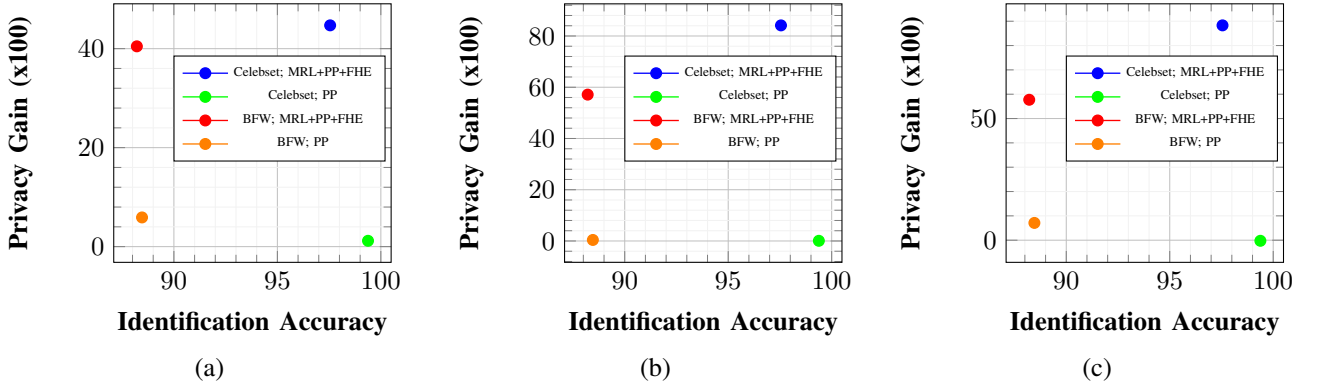
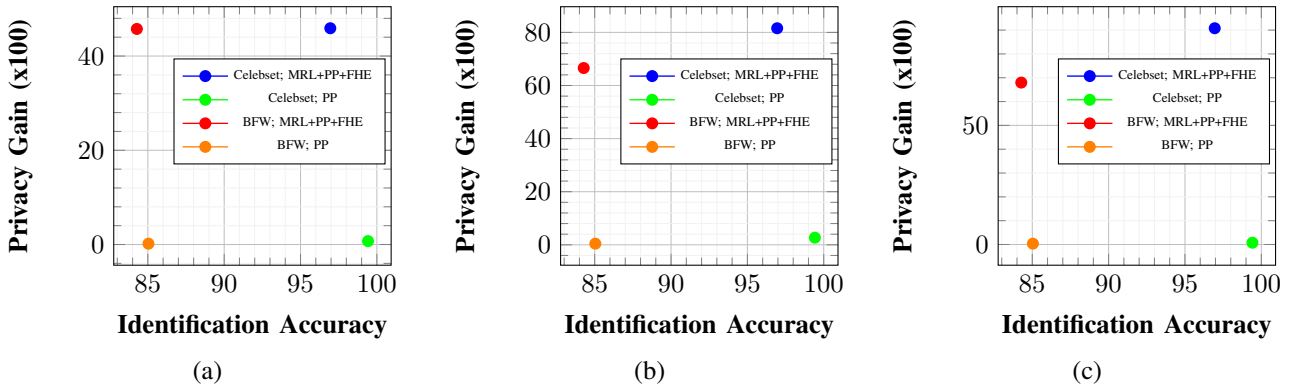
Additionally, we find that utilizing a soft biometrics classifier network trained on plaintext data allows for seamless inference within the FHE domain, yielding the same predictive performance as that in the unencrypted domain (as depicted in Table III in "None"). This underscores the feasibility of conducting not only identification but also soft biometric analysis in the secure FHE domain.

## IX. CONCLUSION

In this paper, we propose to use FHE in combination with template protection and compression to secure the face template and prevent soft biometric leakage. We show that soft biometric attributes from face embeddings can be strictly protected while preserving identification accuracy. In our approach, we compress the face embeddings using MRL (Matryoksha Representation Learning), encrypt them, and then apply PolyProtect as the template protection scheme. The identification performance of the encrypted template compared with the unencrypted version is unchanged. Since FHE guarantees are based on strong theoretical principles,

TABLE V: Suppression Rate across different soft biometric attributes.

Embeddings	Dataset	Template Protection	Identification Accuracy	Gender Attribute	Age Attribute	Ethnicity Attribute
FaceNet	CelebSet	PolyProtect	99.42%	0.0073	0.0306	0.0075
		MRL	97.93%	0.0012	0.0206	0.0145
		MRL + PolyProtect	97.00%	0.0183	0.0041	0.0037
		MRL + FHE	97.83%	<b>0.4678</b>	<b>0.9302</b>	<b>0.9189</b>
		MRL + PolyProtect + FHE	96.95%	<b>0.4678</b>	<b>0.9302</b>	0.9187
	BFW	PolyProtect	85.04%	0.0019	0.0046	0.0038
		MRL	84.42%	0.0339	0.0044	0.0035
		MRL + PolyProtect	84.31%	0.0551	0.0689	0.0464
		MRL + FHE	84.38%	0.4753	0.7034	0.7389
AdaFace	CelebSet	PolyProtect	99.38%	0.0122	0.0008	-0.0026
		MRL	97.95%	-0.0072	0.0526	-0.0273
		MRL + PolyProtect	97.59%	0.0011	0.0124	-0.0201
		MRL + FHE	97.91%	<b>0.4613</b>	0.9322	0.9167
		MRL + PolyProtect + FHE	97.55%	<b>0.4613</b>	<b>0.9323</b>	<b>0.9168</b>
	BFW	PolyProtect	88.46%	0.0656	0.40	0.0869
		MRL	88.33%	0.0269	0.0143	0.0647
		MRL + PolyProtect	88.29%	0.0488	0.0159	0.0751
		MRL + FHE	88.23%	0.4444	0.6710	0.7061
		MRL + PolyProtect + FHE	88.21%	<b>0.4498</b>	<b>0.6718</b>	<b>0.7063</b>

Fig. 7: Privacy Gain of our proposed approach compared to baseline (PP) across different attributes - (a) Gender, (b) Age and (c) Ethnicity - using **Adaface** (MRL - Matryoksha Representation Learning; FHE - Fully Homomorphic Encryption; PP - PolyProtect).Fig. 8: Privacy Gain of our proposed approach compared to baseline (PP) across different attributes - (a) Gender, (b) Age and (c) Ethnicity - using **FaceNet** (MRL - Matryoksha Representation Learning; FHE - Fully Homomorphic Encryption; PP - PolyProtect).

privacy and security are ensured, and only authorized individuals with the secret key will be able to access the results from the FHE computation.

*This material is based upon work supported by the Center for Identification Technology Research and the National Science Foundation under Grant Nos. 1841517 and 1822190.*



## REFERENCES

- [1] Anil K. Jain, Arun Ross, and Karthik Nandakumar. *Introduction to Biometrics: A Textbook*. Springer Science Business Media, 2011.
- [2] Yaniv Taigman, Ming Yang, Marc'Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-level performance in face verification. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [3] Minchul Kim, Anil K. Jain, and Xiaoming Liu. Adaface: Quality adaptive margin for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18750–18759, June 2022.
- [4] Stan Z. Li, Anil K. Jain, and Jiankang Deng, editors. *Handbook of Face Recognition*. Springer Cham, 2024.
- [5] Anil K. Jain, Karthik Nandakumar, and Arun Ross. 50 years of biometric research: Accomplishments, challenges, and opportunities. *Pattern Recognition Letters*, 79(C):80–105, August 2016.
- [6] Antitza Dantcheva, Petros Elia, and Arun Ross. What else does your biometric data reveal? a survey on soft biometrics. *IEEE Transactions on Information Forensics and Security*, 11(3):441–467, 2016.
- [7] Denise Almeida, Konstantin Shmarko, and Elizabeth Lomas. The ethics of facial recognition technologies, surveillance, and accountability in an age of artificial intelligence: a comparative analysis of us, eu, and uk regulatory frameworks. *AI and Ethics*, 2(3):377–387, 2022.
- [8] <https://tinyurl.com/4p2af252>, 2023.
- [9] <https://tinyurl.com/3z7ceksu>, 2023.
- [10] Bryce Westlake, Russell Brewer, Thomas Swearingen, Arun Ross, Stephen Patterson, Dana Michalski, Martyn Hole, Katie Logos, Richard Frank, David Bright, and Erin Afana. Developing automated methods to detect and match face and voice biometrics in child sexual abuse videos. *Trends and Issues in Crime and Criminal Justice*, (648):1–15, 2022.
- [11] Vedrana Krivokuća Hahn and Sébastien Marcel. Towards protecting face embeddings in mobile face verification scenarios. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 4(1):117–134, 2022.
- [12] Baiyu Li and Daniele Micciancio. On the security of homomorphic encryption on approximate numbers. In *Advances in Cryptology – EUROCRYPT: 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17–21, Proceedings, Part I*, page 648–677, Berlin, Heidelberg, 2021. Springer-Verlag.
- [13] Arun Ross, Sudipta Banerjee, and Anurag Chowdhury. Deducing health cues from biometric data. *Computer Vision and Image Understanding*, 221:103438, 2022.
- [14] Ronald L. Rivest, Leonard M. Adleman, and Michael L. Dertouzos. On data banks and privacy homomorphisms. *Foundations of secure computation*, 4:169–180, 1978.
- [15] Kundan Munjal and Rekha Bhatia. A systematic review of homomorphic encryption and its contributions in healthcare industry. *Complex & Intelligent Systems*, 9:1–28, 05 2022.
- [16] Shruthi Gorantala, Rob Springer, and Bryant Gipson. Unlocking the potential of fully homomorphic encryption. *Commun. ACM*, 66(5):72–81, Apr 2023.
- [17] Jung Hee Cheon, Andrey Kim, Miran Kim, and Yongsoo Song. Homomorphic encryption for arithmetic of approximate numbers. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology – ASIACRYPT*, pages 409–437, Cham, 2017. Springer International Publishing.
- [18] Blaž Bortolato, Marija Ivanovska, Peter Rot, Janez Križaj, Philipp Terhöst, Naser Damer, Peter Peer, and Vitomir Štruc. Learning privacy-enhancing face representations through feature disentanglement. In *15th IEEE International Conference on Automatic Face and Gesture Recognition*, pages 495–502, 2020.
- [19] Aythami Morales, Julian Fierrez, Ruben Vera-Rodriguez, and Ruben Tolosana. Sensitivenets: Learning agnostic representations with application to face images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(6):2158–2164, 2021.
- [20] Pietro Melzi, Hatem Otroschi Shahreza, Christian Rathgeb, Ruben Tolosana, Ruben Vera-Rodriguez, Julian Fierrez, Sébastien Marcel, and Christoph Busch. Multi-ive: Privacy enhancement of multiple soft-biometrics in face embeddings. In *IEEE/CVF Winter Conference on Applications of Computer Vision Workshops (WACVW)*, pages 323–331, 2023.
- [21] Zohra Rezgui, Nicola Strisciuglio, and Raymond Veldhuis. Enhancing soft biometric face template privacy with mutual information-based image attacks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) Workshops*, pages 1141–1149, January 2024.
- [22] Yushu Zhang, Tao Wang, Ruoyu Zhao, Wenying Wen, and Youwen Zhu. Rapp: Reversible privacy preservation for various face attributes. *IEEE Transactions on Information Forensics and Security*, 18:3074–3087, 2023.
- [23] Vahid Mirjalili, Sebastian Raschka, and Arun Ross. Privacynet: Semi-adversarial networks for multi-attribute face privacy. *IEEE Transactions on Image Processing*, 29:9400–9412, 2020.
- [24] Vishnu Boddeti. Secure face matching using fully homomorphic encryption. In *IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pages 1–10, 2018.
- [25] Joshua J. Engelsma, Anil K. Jain, and Vishnu Naresh Boddeti. Hers: Homomorphically encrypted representation search. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 4(3):349–360, 2022.
- [26] Arun Kumar Jindal, Imtiyazuddin Shaik, Vasudha Vasudha, Srini-vasa Rao Chalamala, Rajan Ma, and Sachin Lodha. Secure and privacy preserving method for biometric template protection using fully homomorphic encryption. In *IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, pages 1127–1134, 2020.
- [27] Vedrana Krivokuća Hahn and Sébastien Marcel. Biometric template protection for neural-network-based face recognition systems: A survey of methods and evaluation techniques. *IEEE Transactions on Information Forensics and Security*, 18:639–666, 2023.
- [28] Inioluwa Deborah Raji, Timnit Gebru, Margaret Mitchell, Joy Buolamwini, Joonseok Lee, and Emily Denton. Saving face: Investigating the ethical concerns of facial recognition auditing. *CoRR*, abs/2001.00964, 2020.
- [29] Joseph Robinson. Balanced faces in the wild. IEEE Dataport, 2022. <https://dx.doi.org/10.21227/nmsj-df12>.
- [30] Aditya Kusupati, Gantavya Bhatt, Aniket Rege, Matthew Wallingford, Aditya Sinha, Vivek Ramanujan, William Howard-Snyder, Kaifeng Chen, Sham Kakade, Prateek Jain, and Ali Farhadi. Matryoshka representation learning. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 30233–30249. Curran Associates, Inc., 2022.
- [31] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. *CoRR*, abs/1503.03832, 2015.
- [32] Minchul Kim, Anil K. Jain, and Xiaoming Liu. Adaface: Quality adaptive margin for face recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18729–18738, 2022.
- [33] Yong Liu, Xinghua Zhu, Jianzong Wang, and Jing Xiao. A quantitative metric for privacy leakage in federated learning. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3065–3069, 2021.
- [34] Isabel Wagner and David Eckhoff. Technical privacy metrics: a systematic survey. *CoRR*, abs/1512.00327, 2015.
- [35] Blaž Meden, Peter Rot, Philipp Terhöst, Naser Damer, Arjan Kuijper, Walter J. Scheirer, Arun Ross, Peter Peer, and Vitomir Štruc. Privacy-enhancing face biometrics: A comprehensive survey. *IEEE Transactions on Information Forensics and Security*, 16:4147–4183, 2021.