

Multi-IVE: Privacy Enhancement of Multiple Soft-Biometrics in Face Embeddings

Pietro Melzi¹, Hatef Otroushi Shahreza^{2,3}, Christian Rathgeb⁴, Ruben Tolosana¹, Ruben Vera-Rodriguez¹, Julian Fierrez¹, Sébastien Marcel^{2,5}, and Christoph Busch^{4,6}

¹Biometrics and Data Pattern Analytics (BiDA) Lab, Universidad Autonoma de Madrid (UAM), Spain

²Biometrics Security and Privacy Group, Idiap Research Institute, Switzerland

³School of Engineering, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

⁴Biometrics and Security Research Group (da/sec), Hochschule Darmstadt (HDA), Germany

⁵School of Criminal Justice, Université de Lausanne (UNIL), Switzerland

⁶Norwegian Biometrics Laboratory (NBL), Norwegian University of Science and Technology (NTNU), Norway

Abstract

This study focuses on the protection of soft-biometric attributes related to the demographic information of individuals that can be extracted from compact representations of face images, called embeddings. We consider a state-of-the-art technology for soft-biometric privacy enhancement, Incremental Variable Elimination (IVE), and propose Multi-IVE, a new method based on IVE to secure multiple soft-biometric attributes simultaneously. Several aspects of this technology are investigated, proposing different approaches to effectively identify and discard multiple soft-biometric attributes contained in face embeddings. In particular, we consider a domain transformation using Principle Component Analysis (PCA), and apply IVE in the PCA domain.

A complete analysis of the proposed Multi-IVE algorithm is carried out studying the embeddings generated by state-of-the-art face feature extractors, predicting soft-biometric attributes contained within them with multiple machine learning classifiers, and providing a cross-database evaluation. The results obtained show the possibility to simultaneously secure multiple soft-biometric attributes and support the application of embedding domain transformations before addressing the enhancement of soft-biometric privacy.

1. Introduction

Biometric characteristics are biological and behavioural characteristics of individuals from which distinguishing, repeatable biometric features (*i.e.* numbers or labels) can be extracted for the purpose of biometric recognition [9]. The use of biometric characteristics for the recognition of in-

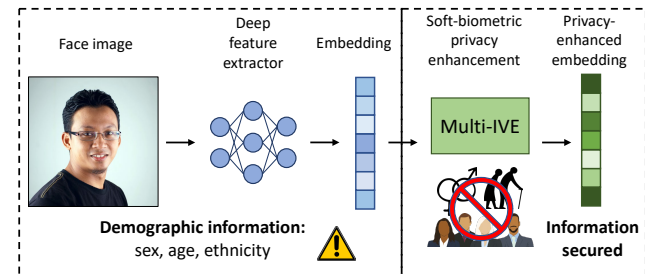


Figure 1: Face image and their embeddings contain sensitive demographic information about individuals, that should be protected. Soft-biometric privacy enhancing technologies secure such information and provide privacy-enhanced embeddings (color image).

dividuals offers considerable advantages compared to traditional knowledge- and possession-based methods (*e.g.* passwords and tokens), as biometric data cannot be forgotten, lost, and transferred to other individuals [1]. Therefore, the use of biometric data is popular and well-accepted in recognition systems, despite privacy concerns about misuse of biometric data [12]. One of these concerns refers to the information that may be extracted from biometric data, including health status, emotions, soft-biometric attributes, and further personal conditions [5]. This information can be obtained without individuals' agreement and used for purposes not initially intended.

This study focuses on the privacy enhancement of soft-biometric attributes that can be found in biometric characteristics. For instance, face images contain many soft-biometric attributes, such as emotions, health status, the color and shape of hair, eyes, the presence of glasses and

other accessories [4]. Soft-biometric attributes that represent demographic information of individuals, *e.g.* sex, age, and ethnicity, are usually found in most of the biometric characteristics (face, iris, fingerprint, gait, voice, etc.) and therefore have received special attention in the literature. Many technologies have been proposed in the literature to secure soft-biometric attributes in biometric data and prevent their disclosure. A description of the general approach followed by these technologies is provided in Figure 1. However, soft-biometric attributes are highly entangled in the representations of biometric data [2], and their protection negatively affects recognition performance.

In this study, we consider a state-of-the-art (SOTA) algorithm for privacy enhancement of soft-biometric attributes. It is called Incremental Variable Elimination (IVE) and was introduced in [18]. It consists of the sequential eliminations of biometric features from the representations of face images, called embeddings, to decrease their contained soft-biometric information. Like other algorithms in the literature [2, 19], IVE was designed to secure a specific soft-biometric attribute at a time. For this reason, we propose a variation of the IVE algorithm to simultaneously secure multiple soft-biometric attributes, in our case sex, age, and ethnicity. If applied *in-cascade*, the same algorithm can secure different soft-biometric attributes contained in embeddings. However, we prefer to secure multiple soft-biometric attributes *simultaneously* for two reasons: *i*) biometric embeddings contain entangled features about multiple soft-biometric attributes [17], and *ii*) *cascaded* applications of the algorithm may remove too much information from embeddings, making them lose their biometric utility.

Furthermore, compared to the original IVE algorithm, we perform a cross-database evaluation of the algorithm and consider embeddings obtained with different SOTA feature extractors for face images. Technologies like IVE require the *strong* (implicit) assumption that soft-biometric information is not equally distributed across embeddings [11]. Hence, we propose transformations of embeddings that precede the application of the IVE algorithm. With such transformations, we investigate how soft-biometric information is contained in embeddings, and aim to satisfy the previous assumption about the distribution of soft-biometric information. We carry out a comprehensive evaluation of Multi-IVE, by assessing how recognition performance and estimation of soft-biometric attributes change during the execution of the algorithm for different experimental settings. To sum up, the main contributions of this study are:

- The proposal of Multi-IVE, a substantial improvement of the original IVE algorithm [18] that allows to secure multiple soft-biometric attributes at the same time. To the best of our knowledge, this is the first technology that simultaneously secures information about sex, age, and ethnicity contained in face embeddings.

- An analysis of the relationship between biometric embeddings and soft-biometric attributes contained therein, considering embedding transformations that address two issues related to the original IVE algorithm: *i*) the distribution of soft-biometric information, and *ii*) the size of secured embeddings.
- An in-depth set of experiments in which, in addition to the aspects already mentioned, we evaluate three variants of the original IVE algorithm to secure multiple soft-biometric attributes, and consider multiple machine learning (ML) classifiers to estimate the soft-biometric attributes contained in face embeddings.
- We make the code available¹ to allow the reproducibility of the experiments presented in this study.

The remainder of the article is organised as follows. In Section 2, we discuss SOTA technologies for soft-biometric privacy enhancement and the original IVE algorithm. In Section 3, we introduce improvements for IVE, namely the embedding transformations and the three variants that consider multiple soft-biometric attributes at the same time. In Section 4, we detail the experimental protocol of the study and specify the databases, SOTA feature extractors, and ML classifiers considered in the study. In Section 5, we provide the results of our experiments and discuss them. Finally, in Section 6, we draw the conclusions of the study.

2. Related works

2.1. Soft-biometric privacy enhancement

To prevent the disclosure of soft-biometric attributes from biometric embeddings, numerous technologies have been proposed in the literature. Their main challenge is to secure soft-biometric attributes without compromising the utility of biometric data for recognition. Most of the technologies listed in the following allow to secure a single soft-biometric attribute at a time. However, it has not been shown how these technologies can be extended and whether *cascaded* applications of them work to secure multiple soft-biometrics. These technologies usually implement the approaches of *data minimisation* or *data protection* to obtain the privacy enhancement of soft-biometric attributes. In data minimisation, the information about soft-biometric attributes is identified in the representations of biometric data and discarded. Consequently, new representations of biometric data are generated excluding soft-biometric information. In data protection, the extraction of soft-biometric attributes is prevented with suitable changes to the original representation of biometric data. In this case, soft-biometric attributes are not discarded, but they are considered inaccessible in the new representations of biometric data [12].

¹<https://github.com/otroshi/multi-ive>

The IVE algorithm described in Section 2.2 implements a data minimisation approach, as specific features are identified and discarded from embeddings. Another technology that follows a similar approach is PFRNet [2]. It is based on an Autoencoder, composed of a two-path encoder and a single-path decoder. The two-path encoder maps face embeddings into two latent vectors: *i*) one that encodes information about identity, and *ii*) the other that encodes information about the sex of individuals. This separation of information into two latent vectors is achieved thanks to different loss functions that disentangle features during the training of the Autoencoder. Therefore, the latent vector that contains soft-biometric information can be discarded.

Contrary to previous approaches implementing data minimisation, other technologies are based on data protection. SensitiveNets is an example of neural networks trained with an adversarial regulariser to learn face representations that protect sex and ethnicity information [14]. PE-MIU is a training-free approach for soft-biometric protection that divides face embeddings into small blocks and randomly changes their positions to make soft-biometric attribute estimation difficult [19]. GaitPrivacyON is a technology developed for mobile gait biometrics, based on training an Autoencoder with different loss functions accounting for gait recognition performance, information about identity, and soft-biometric attributes [6].

2.2. Incremental Variable Elimination (IVE)

The IVE algorithm was presented in [18] to secure a single soft-biometric attribute contained in face embeddings, through the sequential elimination of features from face embeddings. The algorithm is based on the training of a decision tree ensemble to predict a soft-biometric attribute, from which an importance measure for each feature of the embedding can be derived. This measure is used to identify (and eliminate) the features that provide the most significant information about the soft-biometric attribute.

A decision tree is composed of internal and terminal nodes, used to classify embeddings according to some labels. Each internal node represents a binary test that involves a feature of the embedding whereas each terminal node represents the class predicted for the embedding. The importance measure of each feature is derived with a formula that takes into account:

1. the number of embeddings reaching the nodes that represent binary tests involving the feature of interest,
2. the consequent decrease of node impurity, which in turn depends on the proportion of embeddings belonging to each class that traverse the nodes of interest.

At each iteration of the IVE algorithm, a decision tree ensemble is trained to predict a single soft-biometric attribute,

e.g. sex or age, from face embeddings. The importance measure derived from the decision tree ensemble indicates which features to eliminate from embeddings to secure sex or age information. Hence, the IVE algorithm establishes the feature elimination order to secure a soft-biometric attribute in face embeddings. The IVE algorithm works in a scenario of function creep attackers with explicit knowledge of the system privacy mechanisms [12].

3. Proposed method: Multi-IVE

3.1. Transformations of embeddings

The study of the relationship between soft-biometric attributes and biometric embeddings advances the development of soft-biometric privacy enhancing technologies. In [17], a massive attribute classifier was trained to predict 113 (mostly soft-biometrics) attributes contained in deeply-learned face representations, *i.e.* deep embeddings generated by SOTA feature extractors. The study reveals useful information: *i*) the multi-task learning approach used to train the classifier is beneficial, as many attributes share similar features, *ii*) some attributes (*e.g.* sex and age) are strongly correlated with others (*e.g.* beard and accessories), and *iii*) entangled patterns encoded in the embeddings make some attributes easy to predict. As a result, we expect that soft-biometric information related to sex, age, and ethnicity, is spread among numerous features of embeddings.

With the goal of simplifying the distribution of such information, and possibly gathering soft-biometric attributes in a few features, we propose two transformations of the original embeddings that change their domain: *i*) Independent Component Analysis (ICA), and *ii*) Principal Component Analysis (PCA). ICA finds a linear representation of non-Gaussian data made of statistically independent components. Such a representation can capture the structure of data in many applications, including feature extraction and signal separation [8]. Instead, PCA is a linear transformation method that projects data into a space whose orthogonal components retain the maximal variance of data [20]. ICA and PCA provide embeddings in new domains, containing the individual and principal components of the original data, respectively. Given a generic face embedding x we can generate $X_{ICA} = ICA(x)$ and $X_{PCA} = PCA(x)$, such that $X_{ICA} = [ic_1, ic_2, \dots, ic_{n_{ICA}}]$ and $X_{PCA} = [pc_1, pc_2, \dots, pc_{n_{PCA}}]$, with ic_i and pc_i representing the i^{th} components of the new embeddings X_{ICA} and X_{PCA} , $\forall i \in [1, n_{ICA}]$ and $[1, n_{PCA}]$.

Instead of features, Multi-IVE will identify and eliminate components in the transformed domains. Furthermore, the application of ICA and PCA transformations prevents the decrease of embedding size. In fact, embeddings in the transformed domain can be reverted to the original domain (and size). Once reverted, embeddings appear again

in the original domain and lack the information contained in the eliminated components. For instance, by eliminating the principal component pc_2 from X_{PCA} we obtain the embedding $X'_{\text{PCA}} = [pc_1, 0, \dots, pc_{n_{\text{PCA}}}]$ and can revert it to the original domain with the inverse transformation PCA^{-1} , that generates $x' = \text{PCA}^{-1}(X'_{\text{PCA}})$, with $|x| = |x'|$.

3.2. IVE for multiple soft-biometric attributes

As discussed in Section 2.2, the IVE algorithm requires a decision tree ensemble to quantify the importance of each feature in predicting soft-biometric attributes. In this study, we use random forest, as it is a consolidated ensemble learning method and no significant differences with other methods were observed during early experiments. To simultaneously secure multiple soft-biometric attributes, in our case sex, age, and ethnicity, we calculate importance measures for each soft-biometric attribute and combine them according to three variants described in the following. While the first soft-biometric attribute, *i.e.* sex, is binary, we define multiple age intervals and ethnic groups to make the other soft-biometric attributes categorical. We introduce some notations to describe the proposed variants. im_{sb} is the importance measure relative to each soft-biometric attribute $sb \in \{s, a, e\}$, with $s = \text{sex}$, $a = \text{age}$, and $e = \text{ethnicity}$. $x_1, x_2, \dots, x_n = \text{elim}(im)$ are the n features selected according to the importance measure im to be eliminated.

- *Variant A:* At each iteration, we consider one after the other the soft-biometric attributes of interest. For each of them, we calculate the importance measure im_{sb} and eliminate the most important features according to $\text{elim}(im_{sb})$. Hence, at the end of each iteration we have eliminated the features provided by $\text{elim}(im_s) + \text{elim}(im_a) + \text{elim}(im_e)$.
- *Variant B:* We combine the classes c_{sb} of each soft-biometric attribute sb in all possible ways to obtain classes in the format $\{c_s, c_a, c_e\}$. We re-label our data according to the just obtained classes. Then, we apply the original IVE algorithm and eliminate the features provided by $\text{elim}(im_{s \times a \times e})$.
- *Variant C:* We simultaneously calculate importance measures for each soft-biometric attribute, training multiple decision tree ensembles. Then, we sum importance measures, and eliminate the features provided by $\text{elim}(im_s + im_a + im_e)$.

We set the number of features eliminated at each iteration $n = 3$ ($n = 1$ for each sb in variant A) and the maximum number of iterations to 170. In this way, we fairly compare the three variants of the IVE algorithm, as variant A overall eliminates three features, and at most we eliminate 510 features, from face embeddings consisting of 512 features.

4. Experimental settings

4.1. Databases

We use different databases for training and evaluating our Multi-IVE algorithm. For training, we consider the Color FERET database [15], containing 2,722 frontal face images from 994 different individuals with information about sex, age, and ethnicity. We define three age intervals, *i.e.* 0-29, 30-49, 50+, and four ethnicity groups, *i.e.* asian, black, white, and others, according to the labels of the database. We observe that around 63% of images belong to male, 47% to individuals between 30 and 49 years, and 62% to white individuals.

For the evaluation, we consider two databases: *i)* DiveFace, and *ii)* UTKFace. DiveFace has been proposed in [14] to equally represent six classes obtained from the combination of sex with three ethnic groups (asian, black, and white). It contains images from the Megaface dataset [10]. In our study, we select a subset of 6,000 individuals equally representing the six classes to evaluate recognition performance, and estimate sex and ethnicity during the evolution of the Multi-IVE algorithm. UTKFace is a large-scale face database with a long age span, from 0 to 116 years [22]. We select a subset of 6,000 individuals equally representing the three age intervals defined for training, and estimate age during the evolution of the Multi-IVE algorithm.

4.2. Feature extractors

Deep templates are SOTA representations of face images, obtained from Deep Neural Networks (DNNs) with multiple levels of feature extraction. The development of feature extractors for face images has been favored by the design of increasingly effective loss functions, able to generate highly discriminative features [21]. In this study, we consider ArcFace, a feature extractor based on margin loss, a SOTA loss function for deep templates (embeddings in the following) [7]. To provide a benchmark with other SOTA feature extractors, we consider two novel feature extractors: *i)* MagFace [13] and *ii)* ElasticFace [3]. All these feature extractors require images with the size of 112×112 pixels to extract embeddings made of 512 features.

4.3. Soft-biometric classifiers

We use different ML classifiers to estimate the soft-biometric attributes of interest, *i.e.* sex, age, and ethnicity, during the evolution of the Multi-IVE algorithm. We consider Multilayer Perceptron (MLP) with a single hidden layer and three additional ML classifiers to provide more reliable estimates. From the evaluation of the original IVE algorithm [18], we select two classifiers that provide high estimates of soft-biometric attributes: *i)* Support Vector Machine (SVM) with linear kernel and *ii)* logistic regression (LogReg). While MLP can classify non-linearly separable

data, SVM and LogReg are linear classifiers. Finally, we consider the Extra Tree classifier (ET) to also take into account an ensemble method.

4.4. Training

We apply L2-normalisation and standardisation to our embeddings. Eventually, we also apply ICA or PCA domain transformations and generate first embeddings with 512 components, *i.e.* the same size as the original embeddings. Then, for each iteration of the Multi-IVE algorithm, we generate a mask representing the features (or components) to eliminate from embeddings, store it for future cross-database evaluation, and apply it to the training embeddings. Specific masks are generated for each domain transformation (ICA or PCA), IVE variant, and feature extractor considered (ArcFace, MagFace, and ElasticFace).

4.5. Evaluation

We evaluate face recognition performance and estimate soft-biometric attributes from increasingly secured embeddings with the masks generated at each iteration of Multi-IVE. In case of ICA or PCA transformations, we eliminate features in the transformed domain and revert embeddings to the original domain at the end of each iteration. Evaluation is performed every five iterations of Multi-IVE, and repeated ten times with different seeds. In Section 5, we provide the mean and standard deviation values obtained in the ten executions. The main goal is to measure the trend of recognition performance vs. soft-biometric attribute classification. In addition, in the following section we provide some detailed numbers resulting from the evaluation.

Recognition performance is evaluated for the task of verification, with a set of 6,000 individuals provided with at least three images each. Mated comparison scores are obtained from all possible pairs of images of the same individuals whereas non-mated comparison scores are obtained by pairing an image of each individual with ten images of different random individuals. Similarity scores are computed with Euclidean distance, and Equal Error Rate (ERR) is the metric reporting the recognition performance in a single number.

The estimation of soft-biometric attributes is performed with a set of 6,000 individuals, selected to equally represent each soft-biometric attribute of interest. We select one image for each individual, and split images into train and test sets (with proportions of 70% - 30% and soft-biometric stratification). At each iteration the information contained in embeddings (and their size, if the Multi-IVE algorithm is applied in the original domain) will change due to feature elimination. Hence, to estimate each soft-biometric attribute at a given iteration, we have to retrain ML classifiers and evaluate their accuracy. Additionally, every 25 iterations of Multi-IVE we perform a grid search for important

hyper-parameters of ML classifiers, and fine-tune them until the next grid search. These hyper-parameters are: the hidden layer size for MLP, the regularisation parameter for SVM and LogReg, and the number of estimators for ET.

We design four experiments to compare different aspects of Multi-IVE. We present the experiments and the achieved performance in Section 5, in terms of EER for biometric recognition and accuracy for soft-biometric estimation. Even if we estimate in each experiment the three soft-biometric attributes, *i.e.* sex, age, and ethnicity, we report the accuracy obtained for each of them only in the first experiment. To focus on the most important findings, from the second experiment we only report the average accuracy obtained for the three soft-biometrics.

5. Results

5.1. Embedding transformations

In the first experiment, we evaluate Multi-IVE (variant A) for different domain transformations applied to embeddings. Hence, we consider a single feature extractor, *i.e.* ArcFace, and a single soft-biometric classifier, *i.e.* MLP. In addition, we compare Multi-IVE to a baseline algorithm, in which the feature elimination order is random. The performances obtained in the different executions of the algorithms are shown in Figure 2 and described in the following.

Firstly, we compare the Multi-IVE and baseline algorithms in the original domain to assess the advantages provided by the former. For the entire executions of the two algorithms, we observe the same performance in both the tasks of verification and estimation of soft-biometric attributes. This suggests that the feature elimination order established by IVE is no more effective than a random order to secure the soft-biometric information contained in ArcFace embeddings. Moreover, the performance obtained here is comparable to the one provided by the original IVE algorithm in [18]. In both evaluations, recognition performance considerably gets worse when about 400 features are eliminated (with EER relative increase of ca. 15% between the elimination of 300 and 400 features, and ca. 200% between the elimination of 400 and 500 features), while sex and age accuracies slightly decrease to values of 85% and 60%. Further feature eliminations are required to reduce these accuracies, resulting in further decreases in embedding size and recognition performance. Finally, we observe that until the elimination of 400 features ethnicity is almost as easy to predict as sex, despite it is a categorical attribute.

With the goal of making the Multi-IVE algorithm convenient and avoiding reductions in embedding size, we introduce the ICA and PCA domain transformations. With these transformations, we can eliminate components in the transformed domains and verify if soft-biometric information can be secured in the most effective way. ICA and PCA

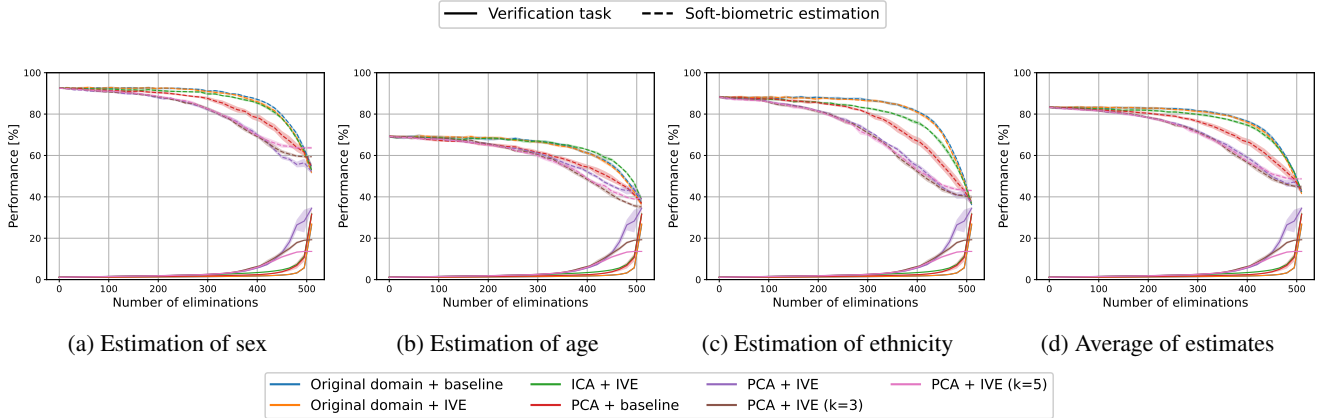


Figure 2: Performance rates obtained for different domain transformations (*i.e.* original, ICA, and PCA) with the Multi-IVE algorithm (variant A) and a baseline algorithm with random feature elimination order. We report performance in terms of EER for the verification task and accuracy for soft-biometric estimation, k represents the number of first principal components locked in PCA domain (color image).

domain transformations have been used in [16] to reduce embedding size, providing embeddings of 128 features able to maintain the recognition performance of original embeddings. However, we generate embeddings in ICA and PCA domains that maintain the same size of the original embeddings to easily compare the different executions of the algorithm. The performance provided by embeddings in the ICA domain is similar to the previous ones, especially for sex and age estimation, with not significant variations observed for EER in the verification task and accuracy in ethnicity estimation.

Differently, the introduction of PCA considerably changes the performance of the algorithm. The situation observed in previous executions when 400 features (or independent components) had been eliminated now appears after the elimination of 300 principal components. Indeed, we get accuracies of 82% for sex, 61% for age, and 72% for ethnicity, and EER increases of 32% between the elimination of 200 and 300 principal components, and 140% between the elimination of 300 and 400 principal components. We also compare the Multi-IVE and baseline algorithms in the PCA domain. This time the two algorithms provide different performances. The baseline algorithm in the PCA domain provides a recognition performance equivalent to those obtained in the original domain and shows lower accuracy in estimating soft-biometric attributes. At the cost of higher EER, the Multi-IVE algorithm in the PCA domain further reduces the accuracy of soft-biometric estimates. Therefore, Multi-IVE can effectively identify the principal components that contain information about soft-biometric attributes in the PCA domain. However, the combination of PCA domain transformation and Multi-IVE leads to rapid increases of EER in the verification task. To face this prob-

lem, we propose to lock the first principal components of PCA embeddings, preventing their elimination during the execution of Multi-IVE. We introduce a parameter k to represent the number of first principal components locked and evaluate its impact on performance. At the end of the algorithm, we observe a small increase in the average accuracy of soft-biometric estimates (from 42.8% to 48.7% with $k = 5$), but a more significant decrease of EER from 34.5% to 13.6% with $k = 5$.

5.2. Feature extractors

As previously discussed, we must consider two important aspects: *i)* the IVE algorithm requires the strong assumption that soft-biometric information is not equally distributed across embeddings [11], and *ii)* soft-biometric attributes share similar features and correlate with other attributes of the embeddings generated by deep feature extractors. Hence, the assumption about the distribution of soft-biometric information is difficult to assess and enforce. In Section 5.1 we observed that the IVE algorithm does not provide any advantage when applied to ArcFace embeddings in the original domain. In this experiment, we evaluate the Multi-IVE algorithm (variant A) with embeddings generated by two other SOTA feature extractors for face images, *i.e.* MagFace [13] and ElasticFace [3]. For each feature extractor we execute: *i)* the baseline algorithm in the original domain, *ii)* the Multi-IVE algorithm in the original domain, and *iii)* the Multi-IVE algorithm in the PCA domain. We report results in Figure 3, accuracy refers to the average estimate of the three soft-biometric attributes obtained with the MLP classifier.

MagFace embeddings provide very similar results to ArcFace embeddings, in both original and PCA domains. In-

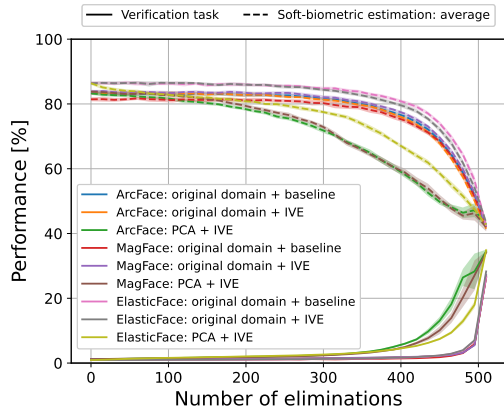


Figure 3: Performance rates obtained for SOTA feature extractors in original and PCA domains. We report performance in terms of EER for the verification task and accuracy for soft-biometric estimation (color image).

terestingly, the accuracy of soft-biometric estimates obtained with Multi-IVE in the original domain is higher than the one obtained with the baseline algorithm, confirming the problems of the IVE algorithm. ElasticFace embeddings provide similar behaviour: no significant difference between the Multi-IVE and baseline algorithms in the original domain, and accuracy decrease with contextual EER increase in the PCA domain. However, it is easier to estimate soft-biometric attributes from ElasticFace embeddings than from other embeddings. According to [17], 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 [3].

5.3. Variants of IVE

We introduced in Section 3.2 three variants of the Multi-IVE algorithm to combine importance measures of different soft-biometric attributes. In this experiment, we evaluate the three variants with ArcFace embeddings transformed in the PCA domain. We consider the following two settings: *i*) no principal component locking, and *ii*) locking of the first five principal components ($k = 5$). We report results in Figure 4, accuracy refers to the average estimate of the three soft-biometric attributes obtained with the MLP classifier.

Variants A and C of Multi-IVE present the same behaviour until the very last iterations, where the effect of $k = 5$ is visible but not the use of different variants. With variant B we observe a faster decrease of accuracy when 200 principal components are eliminated, and a faster increase of EER when 300 principal components are eliminated. Moreover, variant B without principal component




Figure 4: Performance rates obtained for variants of the Multi-IVE algorithm in PCA domain. We report performance in terms of EER for the verification task and accuracy for soft-biometric estimation (color image).

locking provides an unexpected behaviour after the elimination of 400 principal components. For consecutive iterations, despite the usual elimination of principal components, we observe an increase in accuracy and a decrease in EER. This may be due to the fact that the principal components that significantly distinguish the embeddings have already been removed, leaving embeddings with not important or noisy components. This is not the case of $k = 5$ because the first five principal components retain discriminative information in the embeddings. However, all the executions of variant B converge to the same values provided by the respective executions of variants A and C.

5.4. Soft-biometric classifiers

Finally, we compare the estimates of soft-biometric attributes obtained with different ML classifiers from ArcFace embeddings. As in the previous experiment, we consider the three variants of the Multi-IVE algorithm and two settings for the executions of the algorithm in the PCA domain: *i*) no principal component locking, and *ii*) $k = 5$. We report results in Figure 5, accuracy refers to the average estimate of the three soft-biometric attributes. At least until 300 principal components have been eliminated, we obtain the highest estimates of soft-biometric attributes with the MLP classifier. After that, the linear classifiers SVM and LogReg provide better accuracies. Across all the executions, the ET classifier provides the worst estimates.

This experiment highlights the advantage of considering multiple ML classifiers to estimate soft-biometric attributes. In Section 5.1 we provided estimates obtained with the MLP classifier for embeddings secured with the variant A of our Multi-IVE algorithm. Here we observe that such estimates are quite accurate, with the maximum negative

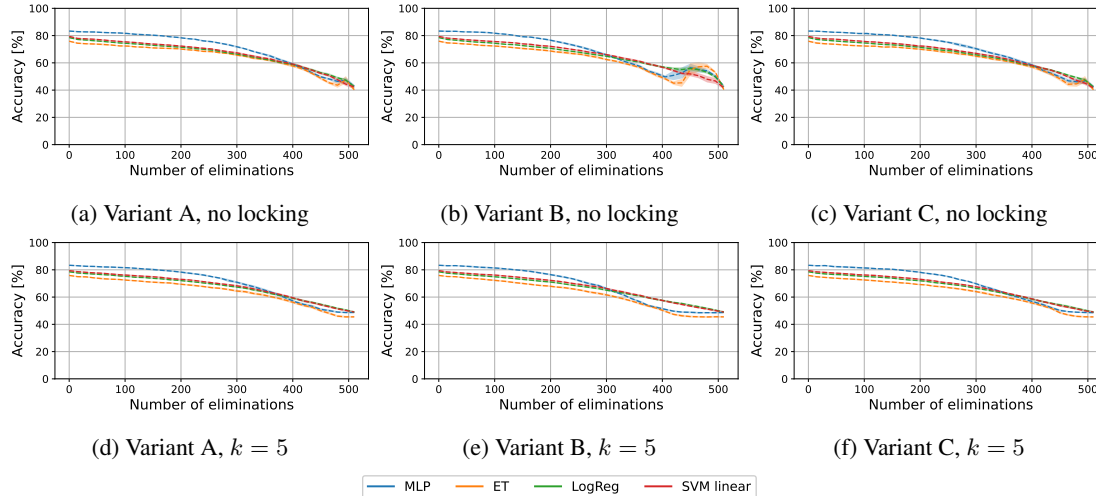


Figure 5: Performance rates obtained for different ML classifiers of soft-biometric attributes in PCA domain. Average estimates of sex, age, and ethnicity are provided. MLP = Multilayer Perceptron, ET = Extra Trees, LogReg = Logistic Regression, SVM = Support Vector Machine (color image).

difference in accuracy between MLP and any other classifier of 3.1% with no principal component locking (Figure 5a) and 3.6% with $k = 5$ (Figure 5d). Also the estimates provided in Section 5.3 for variant C appear quite accurate, with the maximum negative difference in accuracy of 4.1% both in the cases of no principal component locking (Figure 5c) and $k = 5$ (Figure 5f).

However, for variant B we observe higher negative differences in accuracy between MLP and other classifiers of 7% with no principal component locking (Figure 5b) and 6.4% with $k = 5$ (Figure 5e). The estimates provided by SVM and LogReg classifiers for variant B are consistent with those provided by the same classifiers for variants A and C, even if lower than MLP estimates until the elimination of 300 principal components. SVM and LogReg classifiers also avoid the accuracy increase that affects MLP and ET classifiers during the last iterations of the variant B with no principal components locking (Figure 5b). In conclusion, the use of multiple classifiers overall improves the reliability of soft-biometric estimates.

6. Conclusions

In this study, we adapted the IVE algorithm proposed in [18] to the scenario of simultaneous elimination of multiple soft-biometric attributes from face embeddings generated by deep feature extractors. We investigated many aspects of the proposed Multi-IVE algorithm and analysed how soft-biometric information is contained in face embeddings. Domain transformations of embeddings have been proposed when the IVE algorithm struggles to properly secure soft-

biometric information, as in the case of embeddings generated by SOTA feature extractors. In the PCA domain, it is easier to identify the principal components of face embeddings that facilitate the estimation of soft-biometric attributes. However, this study confirms the trade-off between soft-biometric privacy and recognition performance. It can be addressed with specific settings that limit information removal, *e.g.* locking of principal components.

According to the application requirements, we can execute Multi-IVE with the most suitable configuration among those presented in this study, and stop it after a certain number of iterations. It can be the case when we have strong constraints on recognition performance, or want to maintain the non-demographic information of embeddings as much as possible. Future works may define precise configurations of Multi-IVE to target specific applications, and further investigate embedding domain transformation, to effectively identify soft-biometric information in embeddings.

Acknowledgment

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 860813 - TReSPAsS-ETN and the 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. R. Tolosana and R. Vera-Rodriguez are also supported by INTERACTION (PID2021-126521OB-I00 MICINN/FEDER).

References

- [1] Debnath Bhattacharyya and Rahul Ranjan. Biometric authentication: A review. *Science and Technology*, 2(3):16, 2009.
- [2] Blaž Bortolato, Marija Ivanovska, Peter Rot, Janez Križaj, Philipp Terhörst, Naser Damer, Peter Peer, and Vitomir Štruc. Learning privacy-enhancing face representations through feature disentanglement. In *Proc. 2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)*, pages 495–502, 2020.
- [3] Fadi Boutros, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Elasticface: Elastic margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1578–1587, 2022.
- [4] 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.
- [5] Paula Delgado-Santos, Giuseppe Stragapede, Ruben Tolosana, Richard Guest, Farzin Deravi, and Ruben Vera-Rodriguez. A survey of privacy vulnerabilities of mobile device sensors. *ACM Computing Surveys*, 2022.
- [6] Paula Delgado-Santos, Ruben Tolosana, Richard Guest, Ruben Vera-Rodriguez, Farzin Deravi, and Aythami Morales. GaitPrivacyON: Privacy-preserving mobile gait biometrics using unsupervised learning. *Pattern Recognition Letters*, 161:30–37, 2022.
- [7] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019.
- [8] Aapo Hyvärinen and Erkki Oja. Independent component analysis: Algorithms and applications. *Neural networks*, 13(4-5):411–430, 2000.
- [9] ISO/IEC JTC1 SC37 Biometrics. ISO/IEC 2382-37:2022 information technology - vocabulary - part 37: Biometrics, 2022.
- [10] Ira Kemelmacher-Shlizerman, Steven M Seitz, Daniel Miller, and Evan Brossard. The megaface benchmark: 1 million faces for recognition at scale. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4873–4882, 2016.
- [11] Blaž Meden, Peter Rot, Philipp Terhörst, 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*, 2021.
- [12] Pietro Melzi, Christian Rathgeb, Ruben Tolosana, Ruben Vera-Rodriguez, and Christoph Busch. An overview of privacy-enhancing technologies in biometric recognition, 2022. arXiv pre-print.
- [13] Qiang Meng, Shichao Zhao, Zhida Huang, and Feng Zhou. Magface: A universal representation for face recognition and quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14225–14234, 2021.
- [14] 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.
- [15] P Jonathon Phillips, Hyeonjoon Moon, Syed A Rizvi, and Patrick J Rauss. The FERET evaluation methodology for face-recognition algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(10):1090–1104, 2000.
- [16] Philipp Terhörst, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Unsupervised privacy-enhancement of face representations using similarity-sensitive noise transformations. *Applied Intelligence*, 49(8):3043–3060, 2019.
- [17] Philipp Terhörst, Daniel Fährmann, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Beyond identity: What information is stored in biometric face templates? In *2020 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–10. IEEE, 2020.
- [18] Philipp Terhörst, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Suppressing gender and age in face templates using incremental variable elimination. In *Proc. 2019 International Conference on Biometrics (ICB)*, pages 1–8, 2019.
- [19] Philipp Terhörst, Kevin Riehl, Naser Damer, Peter Rot, Blaž Bortolato, Florian Kirchbuchner, Vitomir Štruc, and Arjan Kuijper. PE-MIU: A Training-Free Privacy-Enhancing Face Recognition Approach Based on Minimum Information Units. *IEEE Access*, 8:93635–93647, 2020.
- [20] Michael E Tipping and Christopher M Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622, 1999.
- [21] Mei Wang and Weihong Deng. Deep face recognition: A survey. *Neurocomputing*, 429:215–244, 2021.
- [22] Zhifei Zhang, Yang Song, and Hairong Qi. Age progression/regression by conditional adversarial autoencoder. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2017.