Secure Triplet Loss: Achieving Cancelability and Non-Linkability in End-to-End Deep Biometrics

João Ribeiro Pinto, *Student Member, IEEE,* Miguel V. Correia,
and Jaime S. Cardoso, *Senior Member, IEEE*

Abstract—Biometric systems store sensitive personal data that need to be highly protected. However, state-of-the-art template protection schemes generally consist of separate processes, inspired by salting, hashing, or encryption, that limit the achievable performance. Moreover, these are inadequate to protect current state-of-the-art biometric models as they rely on end-to-end deep learning methods. After proposing the Secure Triplet Loss, focused on template cancelability, we now reformulate it to address the problem of template linkability. Evaluated on biometric verification with off-the-person electrocardiogram (ECG) and unconstrained face images, the proposed method proves successful in training secure biometric models from scratch and adapting a pretrained model to make it secure. The results show that this new formulation of the Secure Triplet Loss succeeds in optimizing end-to-end deep biometric models to verify template cancelability, non-linkability, and non-invertibility.

Index Terms—Biometrics, cancelability, deep learning, invertibility, linkability, security, templates, triplet loss.

1 Introduction

Traditional authentication systems are becoming obsolete as biometric recognition is widely adopted for access control to data and belongings. Biometric systems do not require the user to carry identity cards or remember passwords. Instead, they rely on personal characteristics that are harder to lose, share, or discover than traditional credentials [1], [2], [3].

However, it is easy to change our keys or passwords when a traditional authentication system is compromised, but it is very hard to change our compromised biometric characteristics. This is the reason why it is paramount that biometric templates are kept secure [2], [4]. This is not easily achievable since, unlike password-based systems, biometric comparison is not binary and must also account for the natural intrasubject biometric variability [4], [5].

While several methods have been proposed to protect biometric templates, most require specific feature extraction or additional processes based on salting, biohashing, or cryptographic protection [4], [5]. Even those proposed for deep learning biometric methods [6], [7] are integrated into end-to-end models, thus creating hurdles that often limit the achievable performance.

Considering this, in [8] we proposed the Secure Triplet Loss, a reformulation of the well-known triplet loss that enables training end-to-end deep learning models to obtain cancelable biometric templates. The proposed method allows taking full advantage of the capabilities of end-to-end deep networks while still ensuring the security of the stored biometric data. This methodology was successful in promoting template cancelability and retaining performance levels in ECG biometrics. However, the results have shown that the main drawback of the proposed training loss is failing to promote template non-linkability.

Hence, this paper presents an extension of the aforementioned work focused on tackling this problem. The proposed Secure Triplet Loss is reformulated to include a component that measures and actively promotes template non-linkability. The contributions of this work relative to the prior research in [8] are five-fold:

1) The previously proposed Secure Triplet Loss is reformulated to promote template non-linkability, through a loss component based on the Kullback-Leibler divergence or distance statistics;
2) The evaluation of cancelability is enhanced to test more thoroughly and objectively the proposed methodology and state-of-the-art approaches;
3) Experiments were conducted anew to confirm the solidity of the proposed method for biometric verification, not only with ECG but also with face;
4) The proposed loss formulations are studied in two scenarios: (a) training a model "from scratch" (initialized with random parameters), and (b) adapting an existing end-to-end biometric model to make it secure (taking advantage of pretrained weights and fine-tuning with the proposed method);
5) The Secure Triplet Loss is compared with competitive state-of-the-art approaches based on Bloom Filters [9] and Homomorphic Encryption [10].

As the prior study in [8], this work includes an evaluation of identity verification performance and the template security properties of cancelability, non-linkability, non-invertibility, and secrecy leakage. For realistic results, we use the off-the-person University of Toronto ECG (UofTDB) [11] and the unconstrained YouTube Faces [12] databases, with disjoint sets of identities for training and testing. The proposed Secure Triplet Loss formulations are compared with the original triplet loss in the same evaluation settings.

• *J. R. Pinto, M. V. Correia, and J. S. Cardoso are with INESC TEC and Faculdade de Engenharia da Universidade do Porto, Porto, Portugal.
E-mail: joao.t.pinto@inesctec.pt*

Manuscript received July 24, 2020; revised October 29, 2020.

Besides this introduction, this paper includes the presentation of related concepts and prior art works, in section 2; a detailed description of the original triplet loss and the proposed secure formulations, in section 3; the details on the conducted experiments, in section 4; the results and their discussion, in section 5; and the conclusions drawn from this work, in section 6. Code for this work is available online[1].

2 Background and Related Work

Beyond accounting for natural biometric characteristic variability, biometric data protection methods need to verify template cancelability, non-invertibility, and non-linkability. Cancelability (or revokability) means the templates can be easily and effectively rendered useless if they become compromised, generally through the change of a personal key that is bound with the template [13], [14].

Non-invertibility requires the transformation from biometric samples to templates to be as close to irreversible as possible. Thus, if the template is compromised, the original biometric sample cannot be reliably recovered or approximated [4], [5]. Finally, template non-linkability means it is difficult to assess if compromised templates from different biometric systems belong to the same identity [9].

One of the first template protection methods was the fuzzy commitment scheme proposed by Juels and Wattenberg [15], using cryptography and error-correcting codes for template cancelability. Later, Teoh *et al*. [16] proposed BioHashing, an adaptation of the hashing process commonly applied to passwords to deal with fingerprint variability. A similar approach has been proposed by Sutcu *et al*. [17].

More recently, Rathgeb *et al*. [18], [19] proposed the Bloom filter approach for alignment-free template cancelability and irreversibility. This approach was later adapted by Gomez-Barrero *et al*. [9], [20] to ensure template non-linkability, and by Drozdowski *et al*. [21] for higher computational efficiency. Raja *et al*. [22] proposed a highly efficient method using neighborhood-preserving manifolds and hashing for biometric template protection in smartphones.

Among cryptography-based methods, homomorphic encryption (HE) approaches are particularly promising as HE allows arithmetic operations on the encrypted domain [23]. This allows the biometric comparison to be fully conducted on the encrypted domain, ensuring data security [10]. Fully HE approaches, that allow for unlimited operations in the encrypted domain, most notably include Gentry's [24], Brakerski's [25], and Fan-Vercauteren's [26] schemes.

HE has been successfully applied for biometric template protection in face [10], [23], [27], signature [28], and even multibiometric recognition [29]. However, with HE the protection of templates remains the responsibility of a separate process that should, ideally, be harmoniously integrated within the feature extraction algorithm.

Using deep learning, Pandey *et al*. [6] proposed a template protection scheme that receives features from a convolutional neural network (CNN), quantizes them, and applies hashing to obtain exact comparison despite the variability. Later, Talreja *et al*. [7] used forward error control (FEC) decoding and hashing to protect biometric features extracted by deep neural networks. While these are applied to deep learning, they still require separate protection and comparison schemes. Hence, they are inadequate for recent state-of-the-art biometric recognition methods, which largely rely on end-to-end deep learning models for significantly improved performance.

Considering this, we recently proposed the Secure Triplet Loss [8], a formulation of the triplet loss [30] that enables learning end-to-end models to bind biometric samples with keys. With this method, biometric templates become easily cancelable, just requiring a key change to be invalidated. Additionally, the method provided near-perfect non-invertibility without a decrease in performance relative to the original triplet loss. However, it presented the major drawback of high template linkability.

Hence, this paper addresses this problem by reformulating the Secure Triplet Loss to include a linkability-measuring term. With this, we aim to obtain a general methodology to train end-to-end biometric models that achieve cancelability, non-linkability, and non-invertibility without additional protection processes. The original triplet loss, the original Secure Triplet Loss formulation, and the reformulated Secure Triplet Loss are presented in section 3.

3 Secure Triplet Loss

3.1 Original triplet loss

The triplet loss [30] has been widely used in deep learning to train networks to accurately determine whether or not two samples belong to the same class [31], [32], [33]. During training, such networks receive three inputs (a triplet), in parallel: one is the anchor (x_A, the reference with identity i_A), the second is the positive sample (x_P, with identity $i_P = i_A$), and the third is the negative sample (x_N, with identity $i_N \neq i_A$). In biometrics, triplets are groups of three biometric samples (images or signals): the anchor and positive inputs correspond to the same individual, unlike the negative input.

For each input, the network will output a representation: *e. g.*, for the anchor, $y_A = f(x_A)$. The three representations are then compared using a measure of distance or dissimilarity $d(y_1, y_2)$, and the network is optimized through the minimization of the triplet loss function:

$$l_{TL} = \max\left[0, \alpha + d(y_A, y_P) - d(y_A, y_N)\right], \tag{1}$$

which leads representations of the same class to be more similar than those of different classes, minimizing $d(y_A, y_P)$ and maximizing $d(y_A, y_N)$. The loss also aims to enforce a minimum margin $\alpha > 0$ between the two distances.

This is a generally successful strategy when training neural networks for biometric verification (assessing if the identities of a biometric template and a biometric query match). However, it does not address the important issue of security in biometrics, especially the topics of cancelability and non-linkability.

3.2 Learning cancelability

The training method proposed in [8] modifies the triplet loss to make the final sample representations cancelable (as

[1] SecureTL code repository. Available on: https://github.com/jtrpinto/SecureTL.

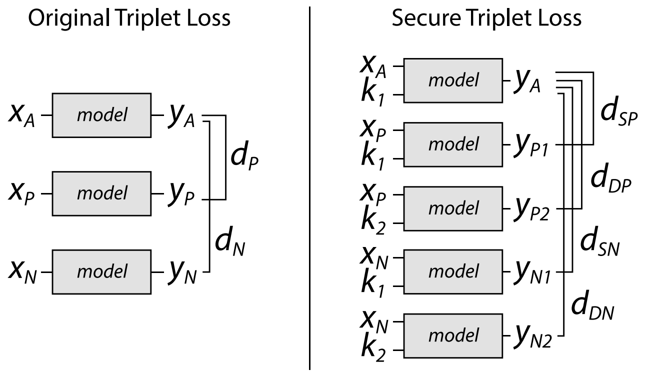

Fig. 1. Comparison between the model training schemes of the original triplet loss and the proposed Secure Triplet Loss method [8].

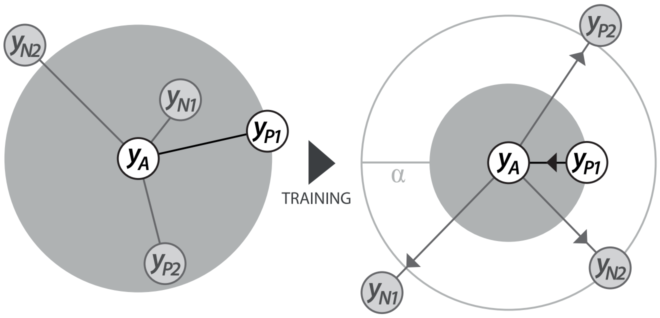

Fig. 2. During training, the Secure Triplet Loss promotes the proximity between y_A and y_{P1} (which match in identity and key) and larger distance to the three negative samples with a margin α.

illustrated in Fig. 1). Besides the triplet inputs (x_A, x_P, and x_N), the network also receives two different keys (k_1, k_2) that are bound with the inputs by the network itself.

Unlike the original triplet loss, x_P and x_N are processed by the network twice. First, they are combined with k_1 and then with k_2. The anchor x_A is only bound with k_1. Thus, five representations are obtained: $y_A = f(x_A, k_1)$, $y_{P1} = f(x_P, k_1)$, $y_{P2} = f(x_P, k_2)$, $y_{N1} = f(x_N, k_1)$, $y_{N2} = f(x_N, k_2)$. From these, four distances are computed: $d_{SP} = d(y_A, y_{P1})$ (with matching identities and keys), $d_{DP} = d(y_A, y_{P2})$ (with matching identities but different keys), $d_{SN} = d(y_A, y_{N1})$ (with different identities but matching keys), and $d_{DN} = d(y_A, y_{N2})$ (with non-matching identities and keys).

The objective is to minimize d_{SP}, when both the identities and the keys match, and maximize the remaining three distances (see Fig. 2). Hence, the loss is computed through:

$$l_{STL} = \max\left(0, \alpha + d_{SP} - d_n\right), \tag{2}$$

where d_n results of the combination of all three distances to be maximized. Here, we consider $d_n = \min(\{d_{SN}, d_{DP}, d_{DN}\})$, with the three distances to be maximized being considered equally relevant. This results in:

$$l_{STL} = \max\left[0, \alpha + d_{SP} - \min(\{d_{SN}, d_{DP}, d_{DN}\})\right]. \tag{3}$$

As with triplet loss, α enforces a margin between positive and negative distances. In this case, the loss involves four distances, since it also takes into account whether or not the keys match. By minimizing the loss in (3), the network learns to deal with the intrasubject and intersubject variability of the biometric characteristic. More importantly, it learns to recognize when the keys do not match, even if the identity is the same. Hence, if the stored templates become compromised, they can easily be invalidated through a key change. However, as reported in [8], l_{STL} fails to promote non-linkability.

3.3 Promoting non-linkability

Non-linkability can be achieved by combining the original formulation of the Secure Triplet Loss, l_{STL}, with a component that quantifies linkability in the representations output by the network during training, l_L. Thus, the proposed reformulation of the Secure Triplet Loss follows the equation:

$$l_{STL2} = \gamma l_{STL} + (1-\gamma) l_L. \tag{4}$$

The l_{STL} component is the original Secure Triplet Loss in (3), focused on biometric performance and template cancelability, following the formulation in (3). The parameter $\gamma \in [0, 1]$ balances the l_{STL} and l_L loss components. The l_L component is focused on measuring template linkability. To achieve non-linkability, one has to ensure similar distance values are obtained when keys don't match (regardless of whether or not the templates are from the same identity). Hence, d_{DP} and d_{DN} should assume similar values. This can be achieved using the Kullback-Leibler divergence (KLD), computed over each batch. This agrees with the reference linkability metric, which is also inspired by the KLD. In this case, this part of the loss becomes:

$$l_L = D_{KL}(P_{d_{DP}} || P_{d_{DN}}) = \sum P_{d_{DP}} \log\left(\frac{P_{d_{DP}}}{P_{d_{DN}}}\right), \tag{5}$$

where $P_{d_{DP}}$ and $P_{d_{DN}}$ are the probability density functions for the distributions of d_{DP} and d_{DN}, respectively. To obtain these distributions and their respective probability density functions, this part of the loss cannot be computed over each triplet, instead being computed over each batch of triplets. For brevity, this formulation of the Secure Triplet Loss with Kullback-Leibler divergence-based linkability is from now on designated as *SecureTL w/KLD* or *STL w/KLD*.

Alternatively, one can avoid estimating these distributions and the computation of the Kullback-Leibler divergence using simple statistics to promote linkability. If we consider $\mu(d_{DP})$ and $\sigma(d_{DP})$ as the mean and standard deviation, respectively, of the distances d_{DP} on a given batch, and likewise $\mu(d_{DN})$ and $\sigma(d_{DN})$ for the distances d_{DN} on the same batch, then we can reformulate

$$l_L = |\mu(d_{DP}) - \mu(d_{DN})| + |\sigma(d_{DP}) - \sigma(d_{DN})|. \tag{6}$$

This should lead the model to offer embeddings that result in similar distance scores when the keys do not match, regardless of whether or not the identities match, thus avoiding template linkability. Throughout the remainder of this paper, for brevity, the formulation of the Secure Triplet Loss with this statistics-based linkability module is designated as *SecureTL w/SL* or *STL w/SL*.

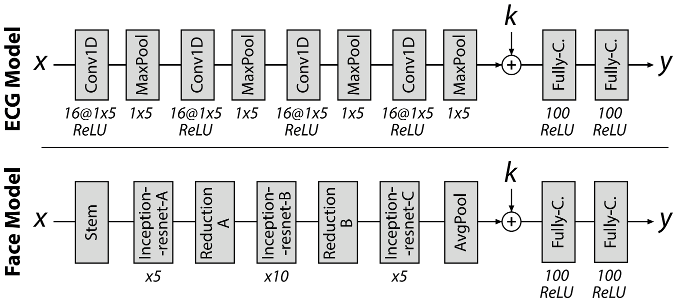

Fig. 3. Architecture of the models used for ECG (top) and face (bottom) identity verification (x denotes a input biometric sample, k a key, and y a biometric template). The structure of the face model before concatenation with k follows exactly the structure of the Inception-ResNet-v1, which is presented in higher detail in [34].

4 Experimental Settings

The proposed methodology for learning secure biometric models was explored for two characteristics: the ECG and face. The ECG comes from the prior research in Secure Triplet Loss [8], and experiments have been conducted anew and in more depth, using PyTorch for more flexibility. The face enables the study of the method's behavior on a more mature and developed biometric characteristic. This section presents the details on the models, the data, and the conducted experiments.

For either characteristic, keys have been randomly generated for each triplet, consisting of unidimensional arrays with 100 binary values. Each key is processed after generation to verify unit-$l2$ norm. For SecureTL w/KLD and SecureTL w/SL, the parameter γ that controls the balance between the Secure Triplet Loss and the linkability component was set to 0.9: this value has overall been able to offer good template non-linkability without considerably harming the validation performance and cancelability.

4.1 ECG identity verification

4.1.1 Data

The ECG data used comes from the University of Toronto ECG Database (UofTDB) [11]. This database includes recordings from 1019 subjects over up to six sessions and five different positions. The signals are off-the-person (less obtrusive and more comfortable for realistic biometric applications) and have been acquired at 200 Hz using dry electrodes on the pointer fingers. Each recording is generally 2 to 5 seconds long.

Data from the last 100 identities were used for training, while the data from the remaining 919 subjects have been reserved for testing. From these 919, one has been discarded for only having a total of 30 seconds of data. Triplets have been generated by selecting an anchor from the first 30 s of data from a subject and positive and negative samples from the remaining data of the same or another identity, respectively. From the 100 training identities, $100\,000$ triplets have been generated, with 20% being used for validation. A total of $10\ 000$ triplets have been generated for testing. Each of the three samples in a triplet is a blindly-segmented five-second raw ECG sample, normalized to zero mean and unit variance.

4.1.2 Model

The model for ECG identity verification (see Fig. 3) follows the architecture of the model used for the previous Secure Triplet Loss research in [8], adapted from the competitive end-to-end model proposed in [31], [35]. The model is composed of four unidimensional convolutional layers (with 16, 16, 32, and 32 filters, respectively, with size 1×5, unit stride, and zero padding), each followed by ReLU activation and max-pooling (with 1×3 kernels and stride 3). The model ends with two fully-connected layers, each with 100 units and followed by ReLU activation.

Once trained, this model receives a 5 second long raw ECG segment (1000 samples long at 200 Hz sampling frequency) and outputs an embedding or template that can be compared to a reference through the Euclidean distance (during training) or through the normalized Euclidean distance [36] (with the trained model, to obtain dissimilarity scores in the $[0, 1]$ range). In the case of Secure Triplet Loss, the feature vector $s(x)$ (the flattened feature maps from the last max-pooling layer) is concatenated with the key array k, and both are bound together by the fully-connected layers to make the final secure template $f(x, k)$.

The model was trained using the Adam [37] optimizer, with initial learning rate of 0.0001 and $l2$ weight regularization with $\lambda = 0.001$. The training lasted a maximum of 250 epochs, with batch size 32, with early stopping based on validation loss with patience of 25 epochs.

4.2 Face identity verification

4.2.1 Data

To fine-tune and evaluate the model, images from the YouTube Faces database [12] were used. This database is composed of frames from 3425 YouTube videos, depicting a total of 1595 subjects (up to six videos of each subject). Each video corresponds to between 48 and 6070 frames. This work used the aligned images provided on the database, which resulted from face detection, cropping, and alignment.

Each face image has been reduced to 70% height and width and resized to 160×160 to match the input dimensions of the model. Ten random triplets have been generated for each of the first 500 subjects on the database for a total of 5000 training triplets, of which 1000 have been used for validation. Ten random triplets have also been generated from each of the remaining identities, reserved for testing, resulting in a total of $10\,950$ test triplets. Whenever possible, the anchor and positive samples corresponded to different videos of the same identity.

4.2.2 Model

The model for face identity verification (see Fig. 3) is based on the Inception-ResNet [34]. This network has been pretrained[2] for identification on the VGGFace2 dataset [38] and offered an accuracy of 99.63% on the Labeled Faces in the Wild (LFW) dataset and 95.12% on the YouTube Faces database [39]. The original fully-connected layer has been replaced with two new fully-connected layers, each with

[2] FaceNet Pytorch Package. Available on: https://github.com/timesler/facenet-pytorch.

100 units and followed by ReLU activation. For the Secure Triplet Loss, the first of these layers receives the feature vector $s(x)$ from the first part of the model, concatenated with the key k. The second outputs the template $y(x, k)$.

All layers on the model have been frozen, to take advantage of the pretrained parameters. The exceptions are the last convolutional block and the fully-connected layers that come, respectively, before and after the average pooling operation. The last convolutional block is fine-tuned to allow for small adjustments during training, while the fully-connected layers are newly created and thus require training. The model was trained for a maximum of 250 epochs at batch size 32, with early stopping based on validation EER with patience of 25 epochs. As with the ECG model, the Adam optimizer was used with an initial learning rate 0.0001 and $l2$ regularization with $\lambda = 0.001$.

4.3 Evaluation frameworks and metrics

The experiments have been designed to quantify the performance of the models trained with the original and Secure Triplet Loss formulations, not only considering verification accuracy, but also biometric security.

4.3.1 Verification performance

The verification performance is quantified through the measurement of false match rates (FMR) and false non-match rates ($FNMR$) over the range of possible decision thresholds (for these models, $t \in [0, 1]$). These values are presented in FMR *vs.* $FNMR$ plots and detection error tradeoff (DET) curves and used to compute the equal error rate (EER), corresponding to the error where $FMR_V = FNMR$, and the $FNMR@FMR = 0.01\%$.

4.3.2 Cancelability

Avoiding additional processes such as biohashing or template encryption, the proposed Secure Triplet Loss integrates cancelability into the single output of the system, the template $y(x, k)$, and is reflected in the distance measure d between two templates. Although the proposed loss is designed to promote cancelability, this property may not necessarily be achieved.

Hence, the experiments with the Secure Triplet Loss include the measurement of cancelability error. The plots of false match *vs.* false non-match rates over the dissimilarity/distance scores include both the false match rate based on identity (when identities don't match, denoted as FMR_V) but also false match rate based on cancelability (when keys don't match, denoted as FMR_C). The false non-match rate ($FNMR$) values are the same for identity and cancelability since they refer to situations when both identity and keys match. The value of cancelability false accept rate at the operation point that corresponds to the verification EER, $FMR_C@EER$, is also computed.

4.3.3 Non-linkability

The template non-linkability analysis followed the method described by Gomez-Barrero *et al*. [9]. The test samples were paired into mated (different biometric samples from the same identity with different keys) and non-mated instances (different identities and keys). These have been used to compute $p(d|H_m)$ and $p(d|H_{nm})$: the probability density functions of the distance score d given the instances are, respectively, mated (hypothesis H_m) or non-mated (hypothesis H_{nm}). From the likelihood ratio $LR(d) = p(d|H_m)/p(d|H_{nm})$, $\mathrm{D}_{\leftrightarrow}(d)$ is computed through

$$D_{\leftrightarrow}(d) = \begin{cases} 0, & \text{if } LR(d) \leq 1 \\ 2\left(\left(1 + \mathrm{e}^{-(LR(d)-1)}\right)^{-1} - \frac{1}{2}\right), & \text{if } LR(d) > 1 \end{cases} \quad (7)$$

which allows to compute the $D_{\leftrightarrow}^{sys}$ linkability metric with

$$D_{\leftrightarrow}^{sys} = \int_{d_{min}}^{d_{max}} D_{\leftrightarrow}(d) \cdot p(d|Hm)\, \mathrm{d}d. \quad (8)$$

The $D_{\leftrightarrow}^{sys}$ is considered the main metric to quantify template linkability. A biometric system verifying perfect template non-linkability, which is highly desirable, will assume $D_{\leftrightarrow}^{sys} = 0$. A biometric system creating entirely linkable templates will verify $D_{\leftrightarrow}^{sys} = 1$.

4.3.4 Non-invertibility and secrecy leakage

Other aspects of template security offered by the proposed method were evaluated, namely non-invertibility and secrecy leakage. Non-invertibility is measured through the privacy leakage rate, which can be computed through the expression:

$$\frac{H(X|Y)}{H(X)} = 1 - \frac{I(X;Y)}{H(X)}, \quad (9)$$

where X is the input biometric, Y is the output of the model, $H(X)$ denotes the entropy of X, $H(X|Y)$ denotes the conditional entropy of X given Y, and $I(X;Y)$ denotes the mutual information between X and Y. The privacy leakage rate, in the range $[0, 1]$, should be as close to 1 as possible: obtaining information on X should be impossible even when one has all knowledge of Y. The secrecy leakage measures the mutual information between the stored template Y and the key K, through the expression $I(Y;K)$. The keys are public, unlike the templates, so they should reveal as little information as possible on the templates. Hence, the secrecy leakage should be close to zero.

These require the computation of some information theoretical measures, such as entropy and mutual information. This is very difficult in biometrics, due to the high dimensionality of the inputs and the feature sets, as well as their variability. In this work, we repeat the process described in [8] to estimate such measures. Entropy and mutual information were estimated using a Python implementation[3] of the methods proposed in [40] and in [41], respectively, for continuous multivariate data. These methods, based on nearest neighbor statistics, were shown to be more accurate than the alternatives [42]. Since the processing cost of such estimations grows exponentially with the size of the dataset, a subset of 1000 test anchors has been used for this test.

5 Results and Discussion

A general overview of the results obtained is presented in Table 1 and Table 2, respectively for ECG and face identity verification. The following subsections discuss the results on verification performance, cancelability, and non-linkability, and the comparison with state-of-the-art alternatives.

[3] Paul Brodersen's Entropy Estimators. Available on: https://github.com/paulbrodersen/entropy_estimators.

TABLE 1
Summary of the test results for ECG identity verification.

Method	Performance		Cancel.	Link.
	EER (%)	$FNMR$ @ $FMR_V = 0.1\%$	FMR_C @ EER	$D_{\leftrightarrow}^{sys}$
Triplet Loss	12.56	0.9033	-	-
STL	11.36	0.8362	0.0035	0.288
STL w/KLD	13.58	0.8700	0.0	0.005
STL w/SL	13.33	0.9458	0.0	0.004
BF [9]	15.76	0.9242	0.0075	0.234
HE [10]	12.49	0.9573	0.0806	0.002

TABLE 2
Summary of the test results for face identity verification.

Method	Performance		Cancel.	Link.
	EER (%)	$FNMR$ @ $FMR_V = 0.1\%$	FMR_C @ EER	$D_{\leftrightarrow}^{sys}$
Triplet Loss	13.99	0.8496	-	-
STL	13.61	0.8314	0.0966	0.399
STL w/KLD	15.93	0.8586	0.0089	0.132
STL w/SL	15.15	0.8771	0.0182	0.070
BF [9]	17.07	0.9103	0.0396	0.245
HE [10]	15.06	0.8312	0.0371	0.001

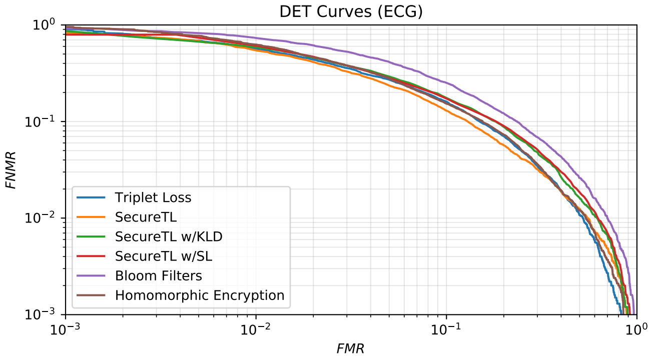

Fig. 4. Detection Error Tradeoff (DET) curves for the ECG identity verification model when trained with the original triplet loss *vs.* the proposed formulations of the Secure Triplet Loss.

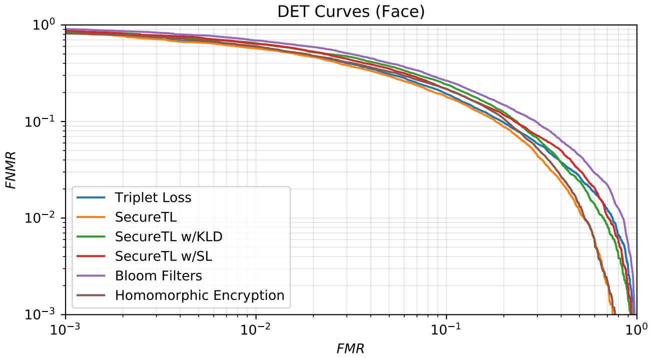

Fig. 5. Detection Error Tradeoff (DET) curves for the face identity verification model when trained with the original triplet loss *vs.* the proposed formulations of the Secure Triplet Loss.

5.1 Verification performance

On ECG identity verification, the baseline method trained with triplet loss offered 12.56% EER. This is similar to the results presented in the work that first proposed this end-to-end model [31], [35]. As presented in Table 1 and in the receiver-operating characteristic (ROC) curves in Fig. 4, the Secure Triplet Loss previously formulated in [8], without considering linkability, attained 11.36% EER, which is an improvement in performance over the original triplet loss despite the inclusion of a cancelability module. Both these results are aligned with those previously reported in the original work in Secure Triplet Loss [8].

The proposed reformulations of the Secure Triplet Loss, which consider template linkability using the Kullback-Leibler divergence (SecureTL w/KLD) or using distance statistics (SecureTL w/SL), led the model to attain, respectively, 13.58% and 13.33% EER. These results show that a small performance gap should be expected when considering both cancelability and linkability in the triplet loss. Recalling the performance improvements with the previous Secure Triplet Loss formulation, it can be hypothesized that the performance decrease in SecureTL w/KLD and SecureTL w/SL is caused by measuring linkability in a separate loss module (computed batch-by-batch and then added to the SecureTL formulation). It is likely that, if linkability was better integrated into the Secure Triplet Loss, as was cancelability, then the performance gap would remain closed.

Nevertheless, the model trained with any of the proposed loss formulations still offers considerably better performance than the state-of-the-art methods. The best state-of-the-art method evaluated in the same conditions (in [31]) offered 21.82% EER *vs.* 13.58% attained by SecureTL w/KLD and 13.33% achieved by SecureTL w/SL. This denotes that the proposed method, while presenting a small performance gap with the linkability loss module, still retains most of the performance advantages associated with deep end-to-end models.

For face identity verification, the performance results are presented in Table 2 and in the ROC curves on Fig. 5. The model trained with the triplet loss attained 13.99% EER, which seems adequate given the difficulty of the evaluation settings: the YouTube Faces provides a challenging framework for evaluation (noted by the 95.12% accuracy achieved by the Inception-ResNet model on this database *vs.* 99.63% on the LFW database), disjoint subsets of identities are used for training/validation and testing, and each identity is only represented by a single template for each comparison (the gallery size is 1).

In harmony with the results on ECG, the model trained with the Secure Triplet Loss without linkability offered a small improvement on verification performance (13.61% EER). Likewise, the addition of a linkability-measuring term to the loss leads to a 2% increase in EER. This confirms the aforementioned belief that the separate linkability loss term is affecting performance and improvements could be achieved by integrating it into the Secure Triplet Loss in a more cohesive way.

Overall, the verification performance results denote that it is possible to adequately train or fine-tune an end-to-end model with the proposed loss formulations. With either biometric characteristic, the performance difference between

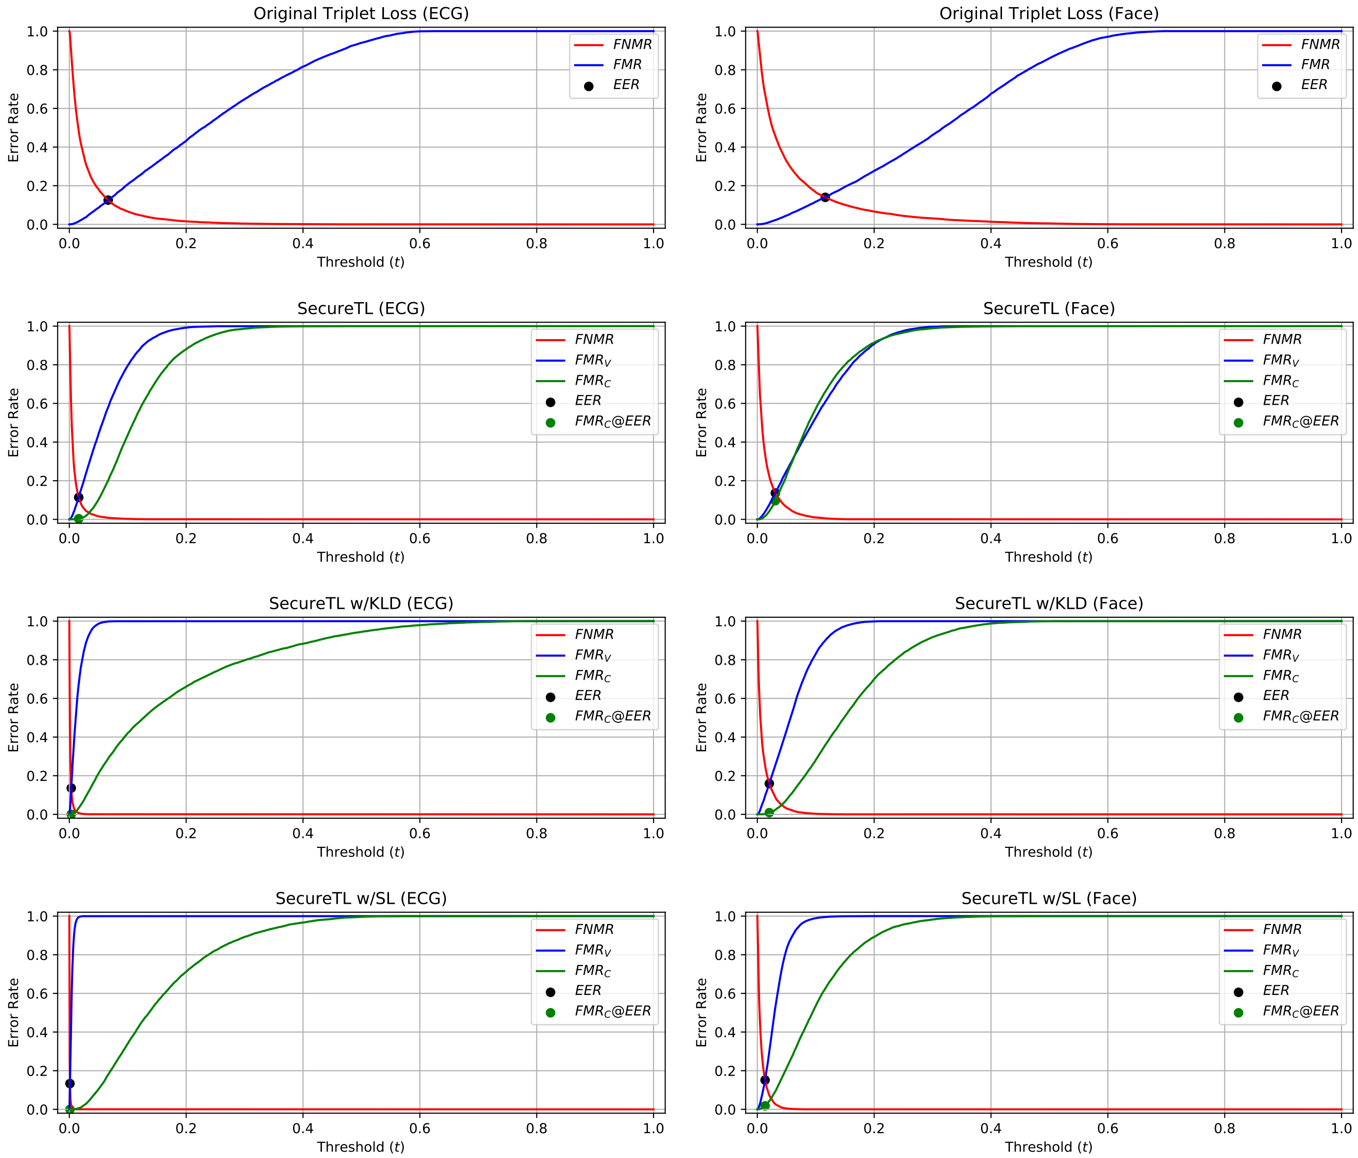

Fig. 6. False match rate (FAR) and false non-match rate ($FNMR$) curves w.r.t. the distance comparison threshold t, for the ECG (left) and face (right) identity verification models with triplet loss and the proposed Secure Triplet Loss formulations. The latter include both FMR_P, relative to verification error, and FMR_C, relative to cancelability error, as well as the FMR_C that corresponds to the EER point ($FMR_C@EER$).

using KLD and distance statistics is not appreciable, which denotes these formulations may each be fitted for specific settings or used interchangeably.

5.2 Cancelability evaluation

As aforementioned, by integrating identity verification and template cancelability into a single comparison score, template cancelability is not necessarily ensured. Hence, the results of false match rates based on cancelability (FMR_C) are presented, in Fig. 6, alongside the false match rates based on verification (FMR_V) and the common false non-match rates ($FNMR$).

In all cases, the FMR_C is lower than FMR_V at and around the EER operation point. In most cases, FMR_C at this point is very small and is lower than or equal to FMR_V for all operation points, which is highly desirable. As presented in Table 1 and Table 2, cancelability error is significantly lower in the ECG models. As shown by the results, SecureTL w/KLD and w/SL appear to be better at promoting cancelability than the original secure loss formulation, denoting that the linkability loss term could have a positive effect on cancelability.

Considering these results and the increased difficulty experienced while fine-tuning the face models, one can conclude that the proposed Secure Triplet Loss is likely better fitted for training models from scratch than to adapt previously trained models to become secure. Nevertheless, the cancelability results, especially with the SecureTL w/KLD and SecureTL w/SL, are encouraging in either case.

5.3 Non-linkability evaluation

The results of the linkability analysis following the framework established in [9] are presented in Fig. 7. In both cases, the original formulation of the Secure Triplet Loss

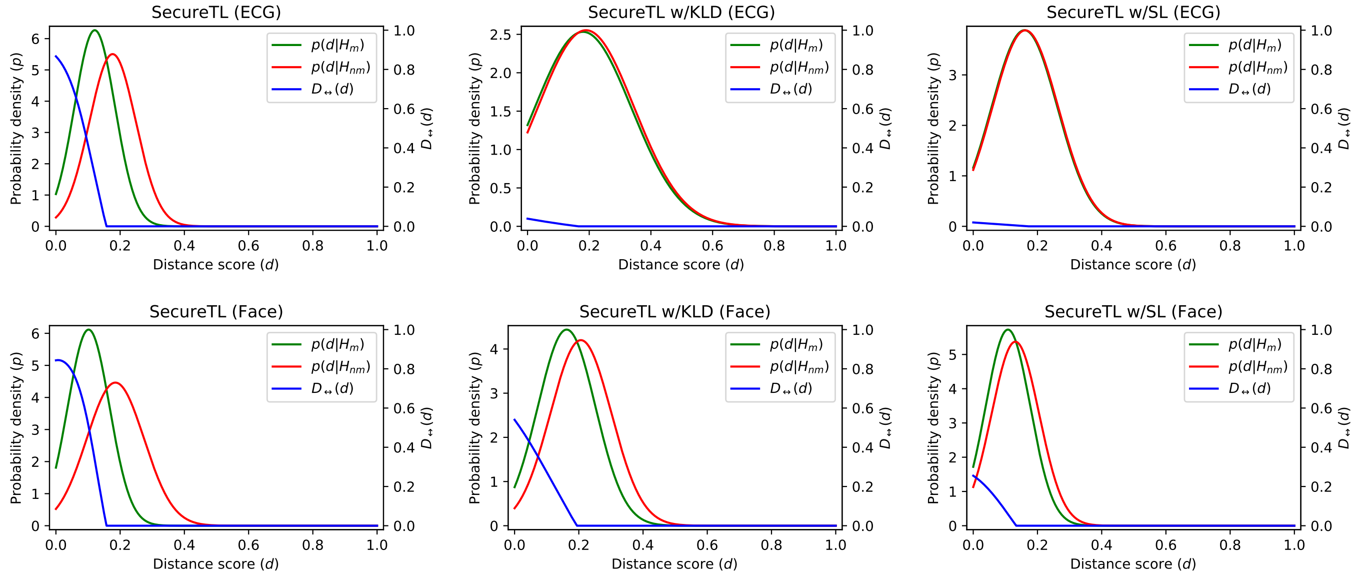

Fig. 7. Template linkability analysis for the ECG (top) and face (bottom) identity verification models, following the procedure proposed in [9].

presents relatively high $D_{\leftrightarrow}^{sys}$ (0.288 for ECG and 0.399 for face). However, the result with ECG is much better than the equivalent reported in the previous work (0.67). This results from the fact that linkability was not promoted by this loss during model training: hence, the model may achieve adequate non-linkability, but that would be accidental.

In the case of the proposed SecureTL w/KLD and SecureTL w/SL, linkability is actively promoted during training through the loss. The effects of this loss reformulation are clear: the probability density functions of mated and non-mated are more superposed, which indicates that it would be more difficult, as desired, to distinguish identities in pairs of templates where the keys do not match.

With ECG, $D_{\leftrightarrow}^{sys}$ assumes the values 0.005 for SecureTL w/KLD and 0.004 for SecureTL w/SL. With face, it assumes 0.132 for SecureTL w/KLD and 0.070 for SecureTL w/SL. All of these can be considered semi to fully-unlinkable. Just as with cancelability, the proposed method seems more adequate for training models from scratch than for fine-tuning existing biometric models. Additionally, using KLD appears to offer some advantages in linkability for ECG verification, but that should be weighted with the increased instability this alternative has shown during training, relative to SecureTL w/SL, especially in face verification.

5.4 Non-invertibility and secrecy leakage

Regarding the non-invertibility and secrecy leakage evaluation, the results follow those previously reported in the original Secure Triplet Loss work [8]. The privacy leakage rate was estimated as 1 for the model trained with any of the losses. This indicates that it is highly difficult for an attacker to recover the original biometric measurements x based on compromised templates y output by the model. As stated in [8], this could be a result of using end-to-end deep learning models: recent research indicates that optimized deep models compress the inputs retaining only the information needed for the task [43]. This means perfect non-invertibility can be achieved without carefully handcrafted feature extraction algorithms.

Similarly, all losses led the model to offer a perfect secrecy leakage rate of 0, which denotes that the public keys used to make the templates cancelable reveal no information on them. These results on non-invertibility and secrecy leakage do not show a superiority of the proposed loss formulations over the original triplet loss but emphasize the meaningful advantages of using end-to-end deep learning models for secure biometrics.

5.5 Comparison with state-of-the-art approaches

The proposed method was compared with two state-of-the-art approaches: Bloom Filters (BF) and Homomorphic Encryption (HE), as described in [9] and [10], respectively. To provide a fair and direct comparison between the template protection schemes, the features given to those methods were those output by the triplet loss baseline model.

The results are presented in Table 1, Table 2, Fig. 4, and Fig. 5. Both with face and ECG, the proposed method outperformed BF in EER, cancelability, and linkability. HE offered the best linkability results, at the cost of poor cancelability. Additionally, HE took significantly longer for biometric comparison than any of the alternatives, which may grant it limited real applicability.

Although the error results are relatively high, the Secure Triplet Loss is competitive *vs.* the state-of-the-art alternatives, especially on cancelability and linkability. Moreover, improved results are expected when the Secure Triplet Loss is used on more accurate biometric models.

5.6 Effects of varying γ

Fig. 8 presents the EER and $D_{\leftrightarrow}^{sys}$ results obtained when varying the γ parameter which balances the original secure triplet loss formulation and the template linkability component. As shown, lower γ values ($\gamma < 0.7$) lead to

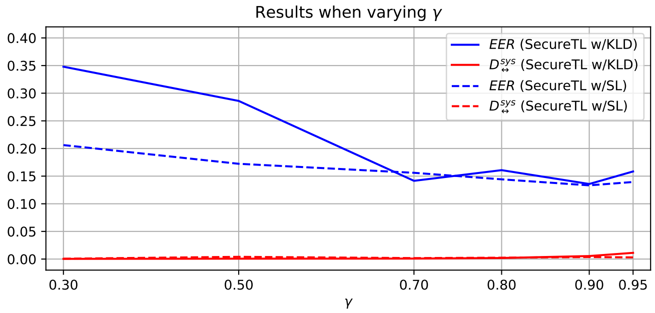

Fig. 8. Results with the proposed loss when varying the γ parameter.

higher EER with either SecureTL w/KLD or SecureTL w/SL, since template non-linkability takes precedence over verification accuracy on the loss that guides model training. For $\gamma \geq 0.7$, lower EER results are obtained, albeit with a slight increase in template linkability ($D_{\leftrightarrow}^{sys}$), especially for $\gamma > 0.9$. Results may vary in other application scenarios depending on their specificities, but $0.7 < \gamma < 0.95$ should offer the highest likelihood of success.

6 Conclusion

This work reformulated the recently proposed Secure Triplet Loss [8] to address the problem of template non-linkability. The goal of this training methodology is to allow the learning of end-to-end deep biometric models, without any additional processes, to verify template cancelability, non-linkability, and non-invertibility. The results on ECG and face identity verification show that the proposed method is not only able to fulfill this purpose, but also to adapt pretrained biometric models to offer secure templates, with competitive performance results.

However, there is still room for improvement. Further efforts should be devoted to design ways to better integrate linkability in the Secure Triplet Loss, in order to avoid performance decreases. A scheme where linkability would be measured triplet-by-triplet (instead of batch-by-batch), similarly to cancelability, should lead to improved performance using the Secure Triplet Loss. This would also enable the formulation of triplet mining approaches for the proposed method. Nevertheless, the Secure Triplet Loss is, overall, a suitable and flexible general scheme for template protection in end-to-end deep biometrics.

Acknowledgments

This work was financed by the ERDF – European Regional Development Fund through the Operational Programme for Competitiveness and Internationalization - COMPETE 2020 Programme and by National Funds through the Portuguese funding agency, FCT – Fundação para a Ciência e a Tecnologia within project "POCI-01-0145-FEDER-030707", and within the Ph.D. grant "SFRH/BD/137720/2018". The authors wish to acknowledge the creators of the UofTDB (University of Toronto, Canada), and the YouTube Faces (Tel Aviv University, Israel) databases, essential for this work.

References

[1] A. K. Jain, A. A. Ross, and K. Nandakumar, *Introduction to Biometrics*. Springer Publishing Company, Incorporated, 2011.

[2] T. Ignatenko and F. M. Willems, "Biometric security from an information-theoretical perspective," *Foundations and Trends in Communications and Information Theory*, vol. 7, no. 2–3, pp. 135–316, 2012.

[3] J. R. Pinto, J. S. Cardoso, and A. Lourenço, "Evolution, Current Challenges, and Future Possibilities in ECG Biometrics," *IEEE Access*, vol. 6, pp. 34 746–34 776, 2018.

[4] K. Nandakumar and A. K. Jain, "Biometric Template Protection: Bridging the performance gap between theory and practice," *IEEE Signal Processing Magazine*, vol. 32, no. 5, pp. 88–100, 2015.

[5] A. K. Jain, K. Nandakumar, and A. Nagar, "Biometric template security," *EURASIP J. Adv. Signal Process*, vol. 2008, pp. 113:1–113:17, 2008.

[6] R. K. Pandey, Y. Zhou, B. U. Kota, and V. Govindaraju, "Deep secure encoding for face template protection," in *2016 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, June 2016, pp. 77–83.

[7] V. Talreja, M. C. Valenti, and N. M. Nasrabadi, "Multibiometric secure system based on deep learning," *2017 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pp. 298–302, 2017.

[8] J. R. Pinto, J. S. Cardoso, and M. V. Correia, "Secure Triplet Loss for End-to-End Deep Biometrics," in *2020 8th International Workshop on Biometrics and Forensics (IWBF)*, Porto, Portugal, Apr 2020.

[9] M. Gomez-Barrero, C. Rathgeb, J. Galbally, C. Busch, and J. Fierrez, "Unlinkable and irreversible biometric template protection based on bloom filters," *Information Sciences*, vol. 370-371, pp. 18–32, 2016.

[10] P. Drozdowski, N. Buchmann, C. Rathgeb, M. Margraf, and C. Busch, "On the Application of Homomorphic Encryption to Face Identification," in *2019 International Conference of the Biometrics Special Interest Group (BIOSIG)*, 2019.

[11] S. Wahabi, S. Pouryayevali, S. Hari, and D. Hatzinakos, "On Evaluating ECG Biometric Systems: Session-Dependence and Body Posture," *IEEE Transactions on Information Forensics and Security*, vol. 9, no. 11, pp. 2002–2013, Nov. 2014.

[12] L. Wolf, T. Hassner, and I. Maoz, "Face recognition in unconstrained videos with matched background similarity," in *CVPR 2011*, June 2011, pp. 529–534.

[13] P. Punithavathi and G. Subbiah, "Can cancellable biometrics preserve privacy?" *Biometric Technology Today*, vol. 2017, no. 7, pp. 8–11, 2017.

[14] M. Tarek, O. Ouda, and T. Hamza, "Robust cancellable biometrics scheme based on neural networks," *IET Biometrics*, vol. 5, no. 3, pp. 220–228, September 2016.

[15] A. Juels and M. Wattenberg, "A Fuzzy Commitment Scheme," in *Proceedings of the 6th ACM Conference on Computer and Communications Security*, New York, NY, USA, 1999, pp. 28–36.

[16] A. Teoh, D. Ngo, and A. Goh, "Biohashing: two factor authentication featuring fingerprint data and tokenised random number," *Pattern Recognition*, vol. 37, no. 11, pp. 2245 – 2255, 2004.

[17] Y. Sutcu, H. T. Sencar, and N. Memon, "A secure biometric authentication scheme based on robust hashing," in *Proceedings of the 7th Workshop on Multimedia and Security*, August 2005, pp. 111–116.

[18] C. Rathgeb, F. Breitinger, and C. Busch, "Alignment-free cancelable iris biometric templates based on adaptive bloom filters," in *2013 International Conference on Biometrics (ICB)*, 2013, pp. 1–8.

[19] C. Rathgeb, M. Gomez-Barrero, C. Busch, J. Galbally, and J. Fierrez, "Towards cancelable multi-biometrics based on bloom filters: a case study on feature level fusion of face and iris," in *3rd International Workshop on Biometrics and Forensics (IWBF 2015)*, 2015, pp. 1–6.

[20] M. Gomez-Barrero, C. Rathgeb, G. Li, R. Ramachandra, J. Galbally, and C. Busch, "Multi-biometric template protection based on bloom filters," *Information Fusion*, vol. 42, pp. 37–50, 2018.

[21] P. Drozdowski, S. Garg, C. Rathgeb, M. Gomez-Barrero, D. Chang, and C. Busch, "Privacy-preserving indexing of iris-codes with cancelable bloom filter-based search structures," in *2018 26th European Signal Processing Conference (EUSIPCO)*, Sep. 2018, pp. 2360–2364.

[22] K. B. Raja, R. Raghavendra, M. Stokkenes, and C. Busch, "Biometric template protection on smartphones using the manifold-structure preserving feature representation," in *Selfie Biometrics: Advances and Challenges*, A. Rattani, R. Derakhshani, and A. Ross, Eds. Cham: Springer International Publishing, 2019, pp. 299–312.

[23] V. N. Boddeti, "Secure face matching using fully homomorphic encryption," in *IEEE International Conference on Biometrics: Theory, Applications, and Systems (BTAS)*, 2018.

[24] C. Gentry, "A fully homomorphic encryption scheme," Ph.D. dissertation, Stanford University, 2009.

[25] Z. Brakerski and V. Vaikuntanathan, "Fully Homomorphic Encryption from Ring-LWE and Security for Key Dependent Messages," in *Advances in Cryptology – CRYPTO 2011*, 2011, pp. 505–524.

[26] J. Fan and F. Vercauteren, "Somewhat practical fully homomorphic encryption," *IACR Cryptology ePrint Archive, Report 2012/144*, 2012.

[27] J. Kolberg, P. Drozdowski, M. Gomez-Barrero, C. Rathgeb, and C. Busch, "Efficiency Analysis of Post-quantum-secure Face Template Protection Schemes based on Homomorphic Encryption," in *2020 International Conference of the Biometrics Special Interest Group (BIOSIG)*, 2020.

[28] M. Gomez-Barrero, J. Fierrez, and J. Galbally, "Variable-length template protection based on homomorphic encryption with application to signature biometrics," in *2016 4th International Conference on Biometrics and Forensics (IWBF)*, 2016.

[29] M. Gomez-Barrero, E. Maiorana, J. Galbally, P. Campisi, and J. Fierrez, "Multi-biometric template protection based on homomorphic encryption," *Pattern Recognition*, vol. 67, pp. 149–163, 2017.

[30] G. Chechik, V. Sharma, U. Shalit, and S. Bengio, "Large scale online learning of image similarity through ranking," *Journal of Machine Learning Research*, vol. 11, pp. 1109–1135, 2010.

[31] J. R. Pinto and J. S. Cardoso, "A end-to-end convolutional neural network for ECG based biometric authentication," in *10th IEEE International Conference on Biometrics: Theory, Applications and Systems (BTAS 2019)*, 2019.

[32] W. Chen, X. Chen, J. Zhang, and K. Huang, "Beyond triplet loss: A deep quadruplet network for person re-identification," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.

[33] D. Cheng, Y. Gong, S. Zhou, J. Wang, and N. Zheng, "Person Re-Identification by Multi-Channel Parts-Based CNN With Improved Triplet Loss Function," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.

[34] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. Alemi, "Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning," in *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)*, 2017.

[35] J. R. Pinto, J. S. Cardoso, and A. Lourenço, "Deep Neural Networks For Biometric Identification Based On Non-Intrusive ECG Acquisitions," in *The Biometric Computing: Recognition and Registration*, K. V. Arya and R. S. Bhadoria, Eds. Boca Raton FL, United States: CRC Press, 2019, ch. 11, pp. 217–234.

[36] Wolfram Language and System Documentation Center, "Normalized square euclidean distance," 2010, (last accessed on 22-11-2019). [Online]. Available: http://reference.wolfram.com/language/ref/NormalizedSquaredEuclideanDistance.html

[37] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in *3rd International Conference for Learning Representations*, 2014.

[38] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "VGGFace2: A dataset for recognising faces across pose and age," in *International Conference on Automatic Face and Gesture Recognition*, 2018.

[39] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Jun 2015.

[40] L. F. Kozachenko and N. N. Leonenko, "Sample estimate of the entropy of a random vector," *Problemy Peredachi Informatsii*, vol. 23, pp. 9–16, 1987.

[41] A. Kraskov, H. Stögbauer, and P. Grassberger, "Estimating mutual information," *Phys. Rev. E*, vol. 69, p. 066138, Jun 2004.

[42] G. Doquire and M. Verleysen, "A comparison of multivariate mutual information estimators for feature selection," in *Proceedings of the 1st International Conference on Pattern Recognition Applications and Methods - Volume 1: ICPRAM,*, 2012, pp. 176–185.

[43] N. Tishby and N. Zaslavsky, "Deep learning and the information bottleneck principle," in *2015 IEEE Information Theory Workshop (ITW)*, April 2015, pp. 1–5.

João Ribeiro Pinto (S'15) is a research assistant with the Visual Computing and Machine Intelligence (VCMI) research group at INESC TEC, Porto, Portugal, and Ph.D. student in Electrical and Computers Engineering at Faculdade de Engenharia da Universidade do Porto (FEUP). In 2017, he has received his M.Sc. degree in Bioengineering (field of Biomedical Engineering) from FEUP, with a thesis on the use of the electrocardiogram (ECG) for biometric recognition of vehicle drivers. Since then, his research has focused on contributing to make the ECG a viable and stronger biometric characteristic in realistic conditions. João's Ph.D. studies are focused on using ECG and face, both acquired almost unnoticeably from vehicle drivers, to recognize them and continuously monitor their wellbeing. His research interests include biometrics, biosignals, pattern recognition, computer vision, and machine learning in general.

Miguel V. Correia graduated in Electrical and Computer Engineering from University of Porto, Faculty of Engineering (FEUP) in 1990. He obtained the Master and the Doctoral degrees also from FEUP in 1995 and 2001, in the fields of Industrial Automation and Computer Vision, respectively. Currently, he is an Assistant Professor at the Department of Electrical and Computer Engineering at FEUP, since 2002 and with tenure since 2007. Since March 2008, he is a senior research member at INESC TEC at Porto, as head of the Bioinstrumentation Lab of the Centre for Biomedical Engineering Research (C-BER). His main research interests are in Electronics and Biomedical Instrumentation, Computational Vision, and Image and Signal Processing, with a focus on sensing methods, technologies, and data fusion for the measurement and analysis of human movement, perception, action, and performance. Since 1990, he participated in more than twenty funded research projects and co-authored over one hundred research papers published in peer-reviewed journals and conference proceedings. He is also a member of the Portuguese Official Engineers Association, the International Association of Pattern Recognition, through its Portuguese chapter, and co-founder of the Portuguese Experimental Psychology Association.

Jaime S. Cardoso (SM'11) holds a Licenciatura (5-year degree) in Electrical and Computer Engineering in 1999, a M.Sc. in Mathematical Engineering in 2005 and a Ph.D. in Computer Vision in 2006, all from the University of Porto. Cardoso is a Full Professor at the Faculty of Engineering, University of Porto (FEUP), where he has been teaching Machine Learning and Computer Vision in Doctoral Programs and multiple courses for graduate studies. Cardoso is currently the coordinator of the Centre for Telecommunications and Multimedia (CTM) at INESC TEC and the co-founder and co-leader of the Breast Research Group and the Visual Computing and Machine Intelligence (VCMI) Group. He is also a Senior Member of IEEE and cofounder of ClusterMedia Labs. His research can be summed up in three major topics: computer vision, machine learning, and decision support systems. Cardoso has co-authored 250+ papers, 80+ of which in international journals. Cardoso has been the recipient of numerous awards, including the Honorable Mention in the Exame Informática Award 2011, in the software category, for the project "Semantic PACS" and First Place in the ICDAR 2013 Music Scores Competition: Staff Removal (task: staff removal with local noise), August 2013. His research results have been recognized both by the peers, with 4800+ citations to his publications, and the advertisement in the mainstream media several times.