

GaitPrivacyON: Privacy-preserving mobile gait biometrics using unsupervised learning

Paula Delgado-Santos^{a,b,*}, Ruben Tolosana^b, Richard Guest^a, Ruben Vera-Rodriguez^b, Farzin Deravi^a, Aythami Morales^b

^a School of Engineering, University of Kent, United Kingdom

^b Biometrics and Data Pattern Analytics Lab, Universidad Autonoma de Madrid, Spain

ARTICLE INFO

Article history:

Received 9 October 2021

Revised 9 May 2022

Accepted 17 July 2022

Available online 20 July 2022

Edited by: Maria De Marsico

Keywords:

Privacy preserving

Sensitive data

Gait verification

Mobile sensors

Biometrics

ABSTRACT

Numerous studies in the literature have already shown the potential of biometrics on mobile devices for authentication purposes. However, it has been shown that, the learning processes associated to biometric systems might expose sensitive personal information about the subjects. This study proposes GaitPrivacyON, a novel mobile gait biometrics verification approach that provides accurate authentication results while preserving the sensitive information of the subject. It comprises two modules: *i*) two convolutional Autoencoders with shared weights that transform attributes of the biometric raw data, such as the gender or the activity being performed, into a new privacy-preserving representation; and *ii*) a mobile gait verification system based on the combination of Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) with a Siamese architecture. The main advantage of GaitPrivacyON is that the first module (convolutional Autoencoders) is trained in an unsupervised way, without specifying the sensitive attributes of the subject to protect. Two experimental studies have been examined: *i*) MotionSense and MobiAct databases; and *ii*) OU-ISIR database. The experimental results achieved suggest the potential of GaitPrivacyON to significantly improve the privacy of the subject while keeping user authentication results higher than 96.6% Area Under the Curve (AUC). To the best of our knowledge, this is the first mobile gait verification approach that considers privacy-preserving methods trained in an unsupervised way.

© 2022 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

1. Introduction

The use of biometrics on mobile devices is currently one of the most popular authentication approaches [1,2]. In particular, behavioural biometrics, which are based on the way subjects perform actions such as writing [3] and walking [4], allow the recognition in a passive way through smart devices, for example, using the accelerometer and gyroscope data [5,6].

Despite the popularity of mobile behavioural biometrics, the data acquired can contain a large amount of personal and sensitive information such as demographics (e.g., gender, age, ethnicity, etc.) or the activity the subject is performing (e.g., walking, sitting, etc.)

* Corresponding author at: School of Engineering, University of Kent, United Kingdom.

E-mail addresses: p.delgado-de-santos@kent.ac.uk (P. Delgado-Santos), ruben.tolosana@uam.es (R. Tolosana), r.m.guest@kent.ac.uk (R. Guest), ruben.vera@uam.es (R. Vera-Rodriguez), f.deravi@kent.ac.uk (F. Deravi), aythami.morales@uam.es (A. Morales).

[7]. As a result, this technology might be considered as an invasion of personal privacy.

Privacy is a concept that has been defined in numerous ways [8], one example of which is the recent General Data Protection Regulation (GDPR) of the European Union [9]. This defines personal data as “any information relating to an identified or identifiable natural person”. Within this set of data, there is a subgroup called sensitive data which includes “racial or ethnic origin, political opinions, religious or philosophical beliefs, trade union membership, genetic data, biometric data for the purpose of uniquely identifying a natural person, health data or data concerning the individual's sex life or sexual orientation”. The automatic processing of such data without the explicit consent of the subject for any specific purpose is prohibited.

The main contributions of this study are:

- A novel mobile gait biometrics verification approach, GaitPrivacyON, that provides accurate authentication results while preserving the privacy of the subject. Fig. 1 represents the general diagram of our proposed approach. It comprises two mod-




Fig. 1. Diagram of GaitPrivacyON, which comprises two modules: *i*) two Autoencoders that are in charge of removing automatically the sensitive data; and *ii*) a gait verification system. Time signals extracted from the accelerometer and gyroscope sensors of the mobile devices are considered as input to GaitPrivacyON. X_e : Enrolled sample, X_t : Test sample, \hat{X}_e : Transformed enrolled sample, \hat{X}_t : Transformed test sample.

ules: *i*) two convolutional Autoencoders with shared weights that transform the biometric raw data into a new privacy-preserving representation (e.g., gender or activity), and *ii*) a mobile gait verification system based on a combination of Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) with a Siamese architecture.

- An in-depth quantitative analysis of GaitPrivacyON over three popular databases in the field of gait recognition, MotionSense [10], MobiAct [11], and OU-ISIR [12], achieving accurate verification results (higher than 96.6% Area Under the Curve, AUC) while reducing the recognition rate of sensitive data to $\sim 50\%$ AUC.
- To the best of our knowledge, this is the first mobile gait verification approach that considers privacy-preserving methods trained in an unsupervised way.

The remainder of the paper is organised as follows. Section 2 summarises previous studies in the field. Section 3 explains all details of our proposed GaitPrivacyON approach. Section 4 summarises the databases considered. Sections 5 and 6 describes the proposed experimental protocol and results, respectively. Finally, Section 7 draws the final conclusions.

2. Related work

2.1. Mobile gait biometrics

Gait biometric recognition allows individuals to be authenticated based on the way they walk. It is a unique characteristic among individuals due to the specific arm swing amplitude, step frequency and length [13]. This characteristic can be easily detected in several ways. One of them is from Inertial Measurement Units (IMU), e.g., accelerometer and gyroscope [13], which enables gait biometrics authentication from mobile devices. An example of this was presented by Mantyjarvi et al. [14]. Gait biometrics data captured by the accelerometer was used in a template matching and cross-correlation framework, achieving together, 7% of Equal Error Rate (EER). Many researchers followed this method, proposing new studies in the literature as described in the review of Sprager and Juric [15].

In recent years, Deep Learning (DL) approaches have dominated the field of gait recognition, being possible to extract more discriminative and robust features. Gadaleta and Rossi created in Gadaleta and Rossi [16] one of the first systems based on DL using CNNs. The authors used universal feature extractors for gait biometrics recognition with misclassification rates of $< 0.15\%$. Their results showed that CNN-based systems learn more useful statisti-

cal features, achieving better performance than previous methods with pre-defined and often arbitrary features.

In addition, RNNs is one of the most powerful DL techniques for temporal sequences [17,18]. Ackerson et al. proposed a new approach in which the OU-ISIR dataset was used [19]. The authors developed one of the first approaches to use a type of RNN, Long Short-Term Memory (LSTM), achieving an EER of 7.55%.

Another interesting approach was proposed by Zou et al. [20]. A hybrid DL model combining CNNs and LSTM for more robust features was created. The proposed model brought together the advances of CNNs (extracting convolutional maps with more discriminative features) and RNNs (processing features as temporal sequences). Mobile devices in the wild were considered, with 118 subjects and data extracted from the accelerometer and gyroscope, obtaining an accuracy of 93.7%.

2.2. Privacy-preserving methods

Privacy-preserving concerns are becoming increasingly important nowadays due to the new privacy laws and regulations. Therefore, many researchers have extensively studied the field in the last decade [8].

In the human-activity recognition field, Iwasawa et al. proposed a model with an adversarial subject classifier and a regular activity-classifier based on CNNs [21]. The authors managed to privatise the subject's discriminative information by 40% while keeping accurate activity recognition performance. Malekzadeh et al. [10] presented a feature learning architecture that provides privacy-preserving data transmission and a new dataset for activity and attribute recognition collected from motion sensors. Their system was based on Generative Adversarial Networks (GANs), achieving a 45.8% reduction in accuracy in the gender classification task while the activity recognition task only decreased by 1.37%. Zhang et al. proposed a new framework for activity recognition and privacy-preserving of sensitive data [22]. The authors wanted to avoid the need for a massive collection of sensitive data for model training. For this purpose, an unsupervised learning training for the privacy-preserving task was performed. The framework was treated by a transformation of the data together with a noise addition consisting of an Autoencoder and a CNN. Results of 56.79% accuracy were achieved for gender classification while the activity recognition task remained almost untouched.

In the gait biometrics verification field, Garofalo et al. developed a Siamese CNN framework [23]. The authors decreased the F1-score in the gender recognition task from 73% to 52% while losing from 90.93% to 85.28% of accuracy in the gait verification task. An adversarial learning technique was used.

Several techniques have also been applied to the image field but at the feature representation level. Therefore, these approaches could also be adapted for time signals. Terhörst et al. proposed an unsupervised approach based on similarity-sensitive noise transformations [24]. That approach added noise to the feature representations. Experiments showed how attackers with prior knowledge about the privacy mechanism (added cosine noise) decreased the accuracy of gender estimation performance with logistic regression $\sim 17\%$. Identity recognition performance only increased by $\sim 5\%$ EER. In [25] Incremental Variable Elimination (IVE) algorithm was proposed. The model was used to suppress binary and categorical attributes in biometric templates. The model, through decision tree training, managed to decrease the gender Correct Overall Classification Rate (COCR) by 20% but only increasing the EER in the identity recognition task by 1.4%. Morales et al. created SensitiveNets [26]. Its main purpose was to set aside sensitive information in decision-making in order to ensure fairness and transparency. It was tested on feature representations of face images by defining and minimising its own loss function. SensitiveNets achieved representations that reduced the gender and ethnicity classification tasks to 54.6% and 53.5% respectively, decreasing recognition task accuracy only 2.6%.

The previous gait verification approaches presented in the literature can preserve specific sensitive attributes [23], but they require a large volume of labelled data for training. On the contrary, GaitPrivacyON considers unsupervised learning for the privacy preserving of the subjects without specifying the sensitive attributes to protect. Therefore, this avoids any model inference and provides greater protection than the models presented in the literature. Also, this approach allows the hiding of all sensitive attributes using a single transformation.

3. Proposed approach: GaitPrivacyON

Fig. 1 shows the general diagram of the proposed privacy-preserving approach. Six time signals are originally acquired from the mobile device as raw data comprising the three axes of the accelerometer and gyroscope. GaitPrivacyON considers a Siamese architecture that is used to learn the similarity between two different biometric templates from the same (genuine) or different (impostor) subject [27]. GaitPrivacyON comprises two modules: *i*) two convolutional Autoencoders with shared weights that transform the biometric raw data into a new privacy-preserving representation (Section 3.1); and *ii*) a mobile gait verification system based on the combination of CNNs and RNNs with a Siamese architecture (Section 3.2). For the training, we adapted the key aspects presented in the image style transformation field [22]. The details are explained in Section 3.3. GaitPrivacyON is an improved adaptation of the approach presented in Zhang et al. [22]. We clarify next the main changes:

- GaitPrivacyON is based on gait biometric verification while the approach presented in Zhang et al. [22] is based on activity recognition. As a result, our approach focuses on verification (1:1) rather than identification (1:N).
- Regarding the Autoencoders considered in GaitPrivacyON (see details in Section 3.1), while TransNet only has a single Autoencoder, two Autoencoders are considered in GaitPrivacyON, sharing their weights through a Siamese architecture. In addition, in order to extract more discriminative features and improve the training, we have considered batch normalization and increased the complexity of the network using more convolutional layers.
- The Gait Verification System proposed in GaitPrivacyON (see details in Section 3.2) has several differences compared with LossNet [22]. LossNet is based only on convolutional layers. The system presented in this work considers both convolutional and

recurrent layers following the state of the art in gait biometrics [20]. This improves the performance of the system and makes our system more robust.

3.1. Autoencoders

Fig. 2 (orange colour) provides a graphical representation of the proposed module. It comprises two convolutional Autoencoders with the same architecture and shared weights. The inputs of each are: enrolled sample (X_e) and test sample (X_t), and the outputs: transformed enrolled sample (\hat{X}_e) and transformed test sample (\hat{X}_t), respectively. The architecture of both is composed of a sequence of 1×3 convolutional filters, coupled with ReLU activation functions. In the encoder, after each convolutional layer, batch normalization and 1×2 max-pooling layers are used to decrease the size of the activation map. In the decoder, after each convolutional layer, a deconvolutional layer is used with 1×3 strides of the convolution. The activation function of the last convolutional layer is linear. In GaitPrivacyON, the loss of the main task (\mathcal{L}_{task}) is in charge of training the Autoencoders to extract useful transformed data (\hat{X}). This loss is considered together with the loss of content ($\mathcal{L}_{content}$), responsible for retaining authentication information, and the loss of style (\mathcal{L}_{style}), which removes sensitive data by introducing uniform random noise (N_s).

3.2. Gait verification system

Fig. 2 (green colour) provides a graphical representation of the architecture proposed for gait verification (φ). In particular, we have adapted the approach originally presented by Zou et al. [20] to our specific case (privacy-preserving gait verification). It is based on a novel Siamese architecture with two inputs: transformed enrolled sample (\hat{X}_e) and transformed test sample (\hat{X}_t). The inputs are reshaped including one new dimension. Unlike the method proposed in Zou et al. [20], the architecture is composed of a sequence of 1×3 two-dimensional convolutional filters, coupled with ReLU activation functions. After 3 convolutional layers, batch normalization, 1×2 max-pooling, and dropout with a probability of 0.5 are used. A reshaping layer is included to return to the shape of the time domain signals followed by a bi-directional LSTM layer with 50 units. The dense layer has a size of 400 with a sigmoid activation function.

3.3. Training

GaitPrivacyON is trained following the idea proposed in the image style transformation field [28]. One image can be divided into two parts: *i*) the *content*, i.e., what is in the image, and *ii*) the *style*, i.e., how the image is illustrated. In our particular application of gait biometrics verification, the content is the unique information that allows to verify the identity of the subject whereas the style is the sensitive information of the subject that can be considered for other purposes not related to the authentication. This sensitive information may include the person's gender, age, ethnicity, or the activity the subject is performing while using mobile devices [21].

Following this idea, three different loss functions have been considered from the work presented in Zhang et al. [22]: *task loss* (\mathcal{L}_{task}), *content loss* ($\mathcal{L}_{content}$), and *style loss* (\mathcal{L}_{style}).

The *task loss* (\mathcal{L}_{task}) helps the system to maintain its usefulness in the main task of gait verification. We consider a categorical cross-entropy that compares the transformed data (\hat{X}) with the biometric raw data (X). The *task loss* can be defined as:

$$\mathcal{L}_{task}(Y_a, \hat{X}) = -Y_a \log(\varphi(\hat{X})) \quad (1)$$

where Y_a and $\varphi(\hat{X})$ are the label and the predicted probability of the gait verification task, respectively.

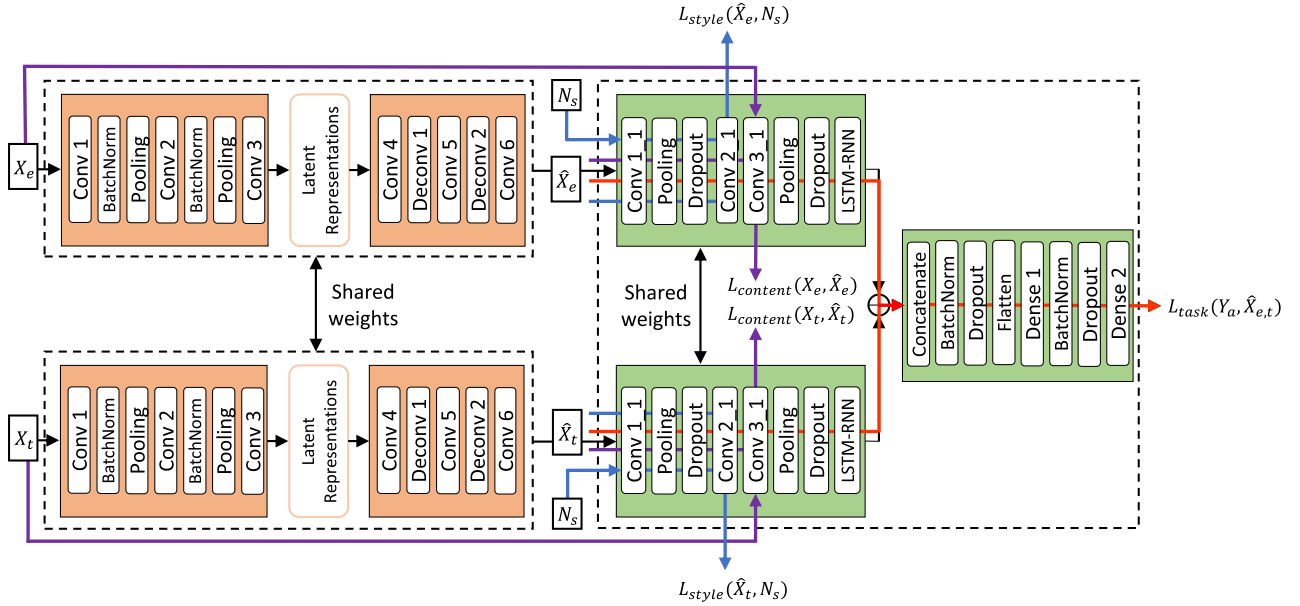


Fig. 2. Architecture and training losses ($\mathcal{L}_{content}$, \mathcal{L}_{style} , \mathcal{L}_{task}) considered in GaitPrivacyON. X_e, X_t : Raw time signals; \hat{X}_e, \hat{X}_t : Transformed time signals; N_s : Random noise; Y_a : label of the gait verification task.

The *content loss* ($\mathcal{L}_{content}$) measures the content (i.e., the authentication information) that the transformed data (\hat{X}) and the biometric raw data (X) have in common. For this aim, we use the Euclidean distance to compare the feature maps provided by the i -layer of the φ network when using both the biometric raw data and the transformed data as input. In our case, we use the feature maps obtained behind *Conv3_1* layer in Fig. 2. This was decided experimentally. The *content loss* is defined as:

$$\mathcal{L}_{content}^i(X, \hat{X}) = \frac{1}{C_i H_i W_i} \|\varphi_i(\hat{X}) - \varphi_i(X)\|_2^2 \quad (2)$$

where i is the layer and $C_i \times H_i \times W_i$ is the shape of the feature map obtained after this layer. Comparing feature maps ensures that the content of the biometric raw data and the transformed data are similar but do not have to be identical.

The *style loss* (\mathcal{L}_{style}) is responsible for maintaining the transformed data (\hat{X}) unstyled, thus avoiding the extraction of any sensitive information automatically. For this purpose, we want to modify the style of the data by uniform random noise (N_s) with range $[-20, 20]$ as done by Zhang et al. [22]. We consider the Gram matrix (G) to measure the style differences between feature representations. Random noise is introduced as the new domain, avoiding using any information from the sensitive data for its protection, creating an unsupervised learning framework. For this aim, both the transformed data and the random noise are fed into the trained gait verification system with the weights frozen. After that, the Gram Matrices of the feature maps obtained as output of the i -layer are compared. The Gram Matrix can be defined as:

$$G_i(X)_{c,c'} = \frac{1}{C_i H_i W_i} \sum_{h=1}^{H_i} \sum_{w=1}^{W_i} \varphi_i(X)_{h,w,c} \varphi_i(X)_{h,w,c'} \quad (3)$$

where the shape of $\varphi_i(X)$ is $C_i \times H_i \times W_i$ and the shape of its Gram matrix (G_i^φ) is $|C_i| \times |C_i|$. $\varphi_i(X)$ can be interpreted as C_i dimensional features for each $H_i \times W_i$ point, where c and c' are two different dimensions.

The *style loss* measures the dissimilarity in style using the Frobenius squared norm of the difference of the Gram matrices of the transformed data (\hat{X}) and the random noise (N_s). In our case, we have decided to use the feature maps obtained behind *Conv2_1*

in Fig. 2. The *style loss* can be defined as:

$$\mathcal{L}_{style}^i(\hat{X}, N_s) = \|G_i^\varphi(\hat{X}) - G_i^\varphi(N_s)\|_F^2 \quad (4)$$

where F denotes the Frobenius squared norm. By using deeper layers, the extracted features will be more similar.

The final loss function of GaitPrivacyON (\mathcal{L}_{total}) would be a weighted sum of the losses \mathcal{L}_{task} , $\mathcal{L}_{content}$, and \mathcal{L}_{style} :

$$\mathcal{L}_{total} = \alpha \mathcal{L}_{task} + \beta \mathcal{L}_{content} + \gamma \mathcal{L}_{style} \quad (5)$$

where $\alpha + \beta + \gamma = 1$.

4. Databases

4.1. MotionSense database

The MotionSense database [10] comprises accelerometer and gyroscope data collected with an iPhone 6 s. A total of 24 subjects with information on gender, age, height, and weight, were obtained. The data was acquired while the subjects performed 4 different activities (walking up and down stairs, jogging, and walking). All the subjects had the mobile phone fixed in the front pocket of the trouser.

4.2. MobiAct database

The MobiAct database [11] comprises accelerometer, gyroscope and magnetometer data collected using a Samsung Galaxy S3. A total of 56 subjects performing the same 4 activities (walking up and down stairs, jogging, and walking) were captured. Data on gender, age, height, and weight of the subjects were acquired. Unlike the previous database, subjects had a free choice of placement of their device, simulating a realistic scenario.

4.3. OU-ISIR database

The popular OU-ISIR database [12] is considered in this study. It comprises accelerometer and gyroscope data collected from three inertial measurement units and a Motorola ME860 around the waist of the subject. In total, 744 subjects with gender and age data were captured. All subjects performed 4 activities (two flat walking, slope-up walking, and slope-down walking).

Table 1

Architecture of the gender and activity inference systems. Prob: Probability; m: number of signals; SAC: Sensitive Attribute Classes.

Layer	Input size ($H \times W \times F$)	Kernel ($H \times W$)	Padding	Activation	Prob
Conv1_1	$m \times 100 \times 1$	1×3	Valid	ReLU	–
Conv1_2	$m \times 98 \times 16$	1×3	Valid	ReLU	–
Batch_1	$m \times 96 \times 16$	–	–	–	–
Pool_1	$m \times 96 \times 16$	1×2	Valid	–	–
Drop_1	$m \times 48 \times 16$	–	–	–	0.5
Conv2_1	$m \times 48 \times 16$	1×5	Valid	ReLU	–
Batch_2	$m \times 44 \times 32$	–	–	–	–
Pool_2	$m \times 22 \times 32$	1×2	Valid	–	–
Drop_2	$m \times 22 \times 32$	–	–	–	0.5
Dense_1	$m \times 100$	–	–	–	–
Batch_3	$m \times 100$	–	–	–	–
Drop_3	$m \times 100$	–	–	–	0.5
Dense_2	$m \times \text{SAC}$	–	–	–	–

5. Experimental protocol

GaitPrivacyON considers two main tasks: *i*) gait biometrics verification, and *ii*) privacy-preserving information, for which auxiliary machine learning systems must be implemented to detect the subject sensitive information, in our case, the gender and activity of the subject while using the mobile device. The specific details of the architecture are included in Section 5.1.

Regarding the training procedure, GaitPrivacyON first trains only the gait verification system using the biometric raw data (X) from the development dataset. For this first stage, binary cross-entropy is considered for the loss function. After that, we train our proposed GaitPrivacyON approach (only the Autoencoders module, the weights of the gait verification system are frozen) using the same development dataset. In this second stage, the total loss function (L_{total}) considered in GaitPrivacyON is a weighted sum of the losses L_{task} , $L_{content}$, and L_{style} , as described in Section 3. The specific details of the development and final evaluation datasets are provided in Section 5.2 and Section 5.3.

5.1. Gender and activity inference systems

Table 1 shows the architecture of the proposed gender and activity inference systems. Six time signals are originally acquired from the mobile device as raw data, the three axes of the accelerometer and gyroscope. The input data is in the same shape as in GaitPrivacyON. The architecture is composed of a sequence of 1×3 convolutional filters, coupled with ReLU activation functions. After some convolutional layers, batch normalization, 1×2 max-pooling, and dropout with a probability of 0.5 are used. The dense layer has a size of 100. For the gender recognition system, a sigmoid activation function is considered whereas softmax is considered for the activity recognition system. Finally, cross-entropy is used for the loss function.

5.2. MotionSense & MobiAct databases

Our approach is trained with accelerometer and gyroscope time signals from both MotionSense and MobiAct databases. A total of 80 subjects (i.e., 24 from MotionSense and 56 from MobiAct) performing 4 different activities (walking up and down stairs, jogging, and walking) are considered in the experimental framework. The total database consists of 55 males and 25 females. In both databases the frequency sampling has been normalised to mean 0 and standard deviation 1, with a sampling frequency of 50 Hz. Each time signal comprises 100 samples. Also, we consider time windows of 2 s with an overlapping ratio of 75%. The total database is divided into development and evaluation datasets, which contain different subjects with random selection. The development dataset, used for the training of GaitPrivacyON, has 70 subjects (85% of the

subjects have been used for training and the remaining part for validation). After training, the remaining 10 unseen subjects are used for the final evaluation. Regarding the gender and activity inference systems (see Section 5.1), we consider the same development and evaluation datasets described before, balancing the number of male and female subjects to avoid bias (5 males and 5 females in the final evaluation set). All subjects contain the same 4 activities.

5.3. OU-ISIR database

GaitPrivacyON is trained with accelerometer and gyroscope time signals using the right-position inertial measurement unit, as it is more reliable according to Ngo et al. [12]. In the scenario of performing 4 different activities (two flat walking, slope-up walking, and slope-down walking), there are 492 subjects available (256 males and 236 females). The data have been normalised with mean 0 and standard deviation 1, with a sampling frequency of 100 Hz. Each time signal has a time window of 1 s, which is defined as 100 samples, and an overlapping between time windows of 75%. This database is divided into development and evaluation, which comprises different subjects with random selection. For the training of GaitPrivacyON, the development dataset contains 80% of the subjects (312 for training and 80 for validation). After the training, the remaining 20% of the subjects (100 unseen subjects) are used for the final evaluation. Regarding the gender and activity inference systems (described in Section 5.1), we consider the same development and evaluation datasets described before, balancing the number of male and female subjects to avoid bias (50 males and 50 females in the final evaluation set). All subjects contain the same 4 activities.

6. Experimental results

6.1. Gender and activity inference from biometric raw data

In this first experiment we analyse the ability of machine learning systems to infer sensitive information of the user from the biometric raw data.

6.1.1. MotionSense & MobiAct databases

Fig. 3 (top) shows the Receiver Operating Characteristic (ROC) curve together with the AUC of the activity recognition system (solid curve). The proposed system achieves 99.2% AUC, differentiating the activity (walking up and down stairs, jogging, and walking) with precision.

Second, we analyse the results achieved by the proposed gender recognition system. The system has two classes: male and female. Fig. 3 (middle) shows the ROC curve together with the AUC achieved by the gender recognition system (solid curve). As in the




Fig. 3. ROC curves and AUC (%) results on the MotionSense and MobiAct evaluation dataset for the two scenarios considered: i) Biometric raw data (X), and ii) GaitPrivacyON (\hat{X}). Different parameters (α, β, γ) of GaitPrivacyON are tested in order to evaluate the results of the main task (gait verification) and the privacy-preserving information of the subject (activity and gender recognition).

case of the activity task, the gender recognition system is able to differentiate the gender with 99.0% AUC.

6.1.2. OU-ISIR database

Fig. 4 (top) shows the ROC curve together with the AUC result achieved by the activity recognition system (solid curve). Similar to the MotionSense and MobiAct databases, the system is able to achieve accurate results with 86.0% AUC. Regarding the gender

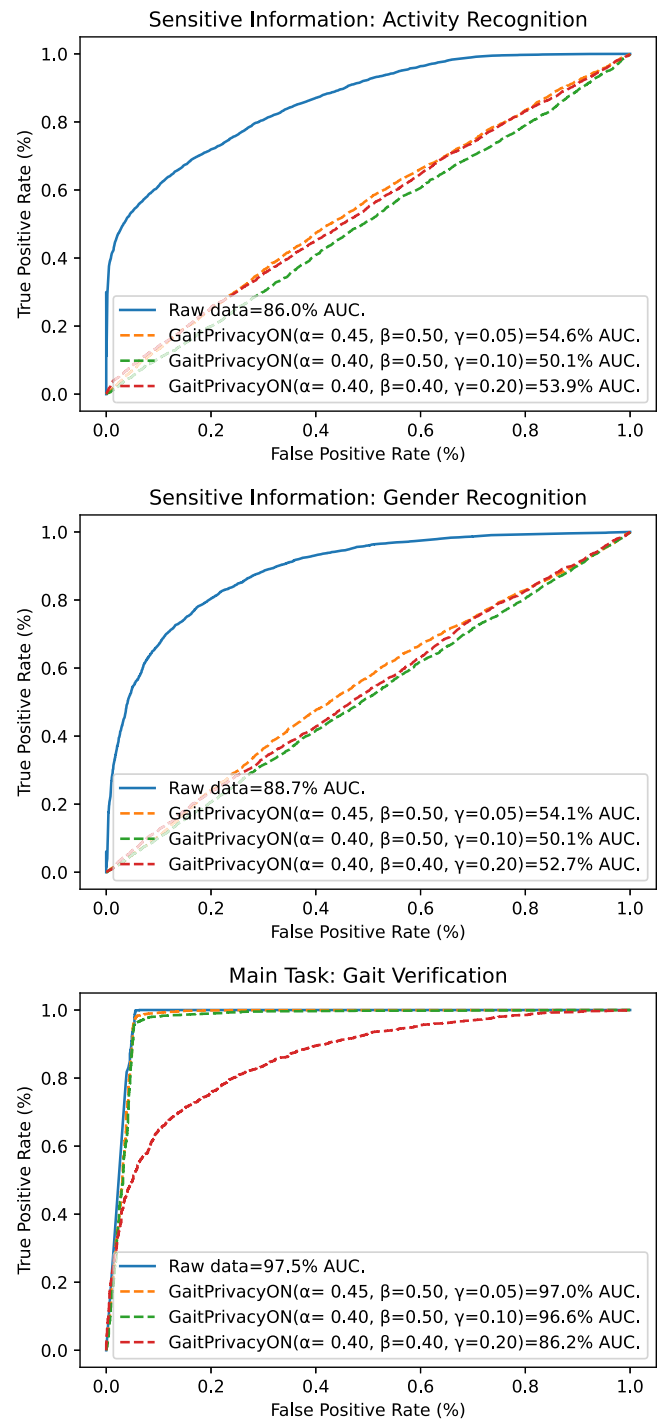


Fig. 4. ROC curves and AUC (%) results on the OU-ISIR evaluation dataset for the two scenarios considered: i) Biometric raw data (X), and ii) GaitPrivacyON (\hat{X}). Different parameters (α, β, γ) of GaitPrivacyON are tested in order to evaluate the results of the main task (gait verification) and the privacy-preserving information of the user (activity and gender recognition).

recognition, see Fig. 4 (middle), good results are also achieved with 88.7% AUC.

These preliminary results support the ability of machine learning systems to infer sensitive information of the subjects from the biometric raw data (X), which might be considered as an invasion of the personal privacy. The next experiments analyse the results achieved by the proposed GaitPrivacyON approach considering the privacy-preserving domain (\hat{X}).

6.2. GaitPrivacyON

As indicated in Section 3.3, three different parameters can be configured in the training process of GaitPrivacyON to control the data transformation and the trade-off between the utility of the gait verification (main task) system and the sensitive information of the user (activity and gender): α (task loss parameter), β (content loss parameter), and γ (style loss parameter).

6.2.1. MotionSense & MobiAct databases

We first analyse the results achieved in the main task, mobile gait verification. Fig. 3 (bottom) shows the ROC curves together with the AUC results of the gait biometrics verification system. Analysing the traditional approach, i.e., using the biometric raw data (X), the gait verification system is able to achieve accurate results with 99.9% AUC over the final evaluation dataset. However, as it was commented in Section 5.1, from this traditional approach it is also possible to extract sensitive user information, 99.2% AUC for activity recognition and 99.0% AUC for gender recognition.

The results achieved by GaitPrivacyON in the main task (gait verification) can be seen in Fig. 3 (bottom). In general, we can see different AUC results depending on the values of the training parameters (including symbols), ranging from 99.6% AUC to 91.0% AUC. The selection of these parameters affects in the extraction of the activity and gender sensitive information.

Fig. 3 (top) shows the ROC curves together with the AUC results achieved by GaitPrivacyON in the activity recognition task (dashed curves) when X is replaced by \hat{X} . It can be seen how the AUC results decrease as γ increases, achieving a result close to random (49.02% AUC) when $\gamma = 0.20$.

A similar trend is also observed for the gender recognition task. Fig. 3 (middle) shows the ROC curves together with the AUC results achieved by GaitPrivacyON in the gender recognition task (dashed curves). It can be seen how the AUC results decrease as γ increases, achieving a result close to random (50.6% AUC) when $\gamma = 0.30$.

As a result, when the transformed data (\hat{X}) provided by GaitPrivacyON achieves AUC values close to random (50.0%) in the sensitive user information tasks, it will be assumed to achieve privacy-preserving results, as long as the AUC of the gait verification task hardly decreases. Therefore, we select as the optimal configuration parameters the $\alpha = 0.40$, $\beta = 0.40$, $\gamma = 0.20$, as the results in the gait biometrics verification task barely decrease (3.15% AUC) while results close to random are achieved in both the activity (49.0% AUC) and gender (57.2% AUC).

6.2.2. OU-ISIR database

Fig. 4 (bottom) shows the ROC curves together with the AUC results of the gait biometrics verification system. Using the biometric raw data (X) of the final evaluation dataset, the gait verification system is able to achieve accurate results with 97.5% AUC. As in the previous case, with this traditional approach it is possible to extract much of the sensitive information such as the activity (86% AUC) and gender (88.7% AUC). For the OU-ISIR database, the best parameter configuration of GaitPrivacyON is $\alpha = 0.40$, $\beta = 0.50$, $\gamma = 0.10$. In this case, GaitPrivacyON achieves AUC results close to 50% for both activity and gender recognition, while keeping a similar performance on gait verification compared with the traditional approach, i.e., 97.5% AUC vs. 96.6% AUC.

6.3. Comparison with the state of the art

Analysing MotionSense and MobiAct databases together, GaitPrivacyON is able to decrease the AUC in the gender task (sensitive information) from 99.0% to 57.2% while reducing the performance from 99.9% AUC to 96.7% AUC in gait verification (main task).

Moreover, using the OU-ISIR database, GaitPrivacyON also achieves robust results, decreasing the AUC from 88.7% to 50.1% in gender recognition while keeping similar AUC results in the main task, from 97.5% to 96.6%. In comparison to our work, the approach presented by Garofalo et al. [23] using the OU-ISIR database decreased the F1-score in the gender recognition task from 73% to 52% while worsening the accuracy from 90.9% to 85.3% in the gait verification task. However, it is important to note that their method considers supervised learning, while GaitPrivacyON is based on unsupervised learning.

Finally, for completeness, we highlight other approaches focused on the privacy-preserving of time sequences [22,29,30], although the topic is different, i.e., activity recognition. A similar trend can be observed when protecting sensitive information such as the age and identity of the person.

7. Conclusions

This study has presented GaitPrivacyON, a novel mobile gait biometrics verification approach that provides accurate authentication results while preserving the privacy of the subject. One of the main advantages of the approach is that the first module (convolutional Autoencoders) is trained in an unsupervised way, without specifying the sensitive attributes of the subject to protect. We have performed an in-depth quantitative analysis of GaitPrivacyON over three popular databases in the field of gait recognition, MotionSense [10], MobiAct [11], and OU-ISIR [12]. Our model is able to obtain good results, as the gait biometrics verification task barely decrease (3.2% AUC with MotionSense and MobiAct databases and 0.9% with OU-ISIR database) while results close to random are achieved in both the activity and gender ($\sim 50\%$ AUC) tasks. In conclusion, GaitPrivacyON increases the protection of the sensitive data (e.g., activity and gender) with unsupervised learning while being able to maintain the accuracy of the gait biometrics verification task. The proposed GaitPrivacyON approach have been evaluated with discrete sensitive attributes (i.e., activity and gender) and further experiments are necessary to adapt the method to continuous sensitive attributes (e.g., weight or age). Our approach is based on a semi-supervised learning approach and therefore, it requires large amount of labelled data (sensitive attributes). Future work will be oriented to: 1) reduce the amount of data needed to train the models using unsupervised approaches; 2) reduce the training time through GPU parallelization.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This project has received funding from the European Unions Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement no 860315. R. Tolosana and R. Vera-Rodriguez are also supported by INTER-ACTION (PID2021-126521OB-I00 MICINN/FEDER).

References

- [1] M. Boakes, R. Guest, F. Deravi, B. Corsetti, Exploring mobile biometric performance through identification of core factors and relationships, *IEEE Trans. Biom., Behav., Identity Sci.* 1 (4) (2019) 278–291.
- [2] E. Maiorana, P. Campisi, N. González-Carballo, A. Neri, Keystroke dynamics authentication for mobile phones, *Proc. ACM Symposium on Applied Computing*(2011).
- [3] L. De Luisa, G.E. Hine, E. Maiorana, P. Campisi, In-air 3D dynamic signature recognition using haptic devices, in: *Proc. IEEE International Workshop on Biometrics and Forensics*, 2021.

- [4] O.C. Reyes, R. Vera-Rodriguez, P. Scully, K.B. Ozanyan, Analysis of spatio-temporal representations for robust footstep recognition with deep residual neural networks, *IEEE Trans. Pattern Anal. Mach. Intell.* 41 (2) (2018) 285–296.
- [5] E. Ellavarason, R. Guest, F. Deravi, R. Sanchez-Riello, B. Corsetti, Touch-dynamics based behavioural biometrics on mobile devices—A review from a usability and performance perspective, *ACM Comput. Surv.* 53 (6) (2020) 1–36.
- [6] A. Acien, A. Morales, R. Vera-Rodriguez, J. Fierrez, R. Tolosana, Multilock: mobile active authentication based on multiple biometric and behavioral patterns, *Proc. International Workshop on Multimodal Understanding and Learning for Embodied Applications*(2019).
- [7] R. Tolosana, J. C. Ruiz-Garcia, R. Vera-Rodriguez, J. Herreros-Rodriguez, S. Romero-Tapiador, A. Morales and J. Fierrez, "Child-computer interaction with mobile devices: recent works, new dataset, and age detection", *IEEE Transactions on Emerging Topics in Computing*, pp. 1-13, 2022.
- [8] P. Delgado-Santos, G. Stragapede, R. Tolosana, R. Guest, F. Deravi, R. Vera-Rodriguez, A survey of privacy vulnerabilities of mobile device sensors, *ACM Comput. Surv.* (2022).
- [9] EU 2016/679 (General Data Protection), 2016, Regulation. <https://gdpr-info.eu/>.
- [10] M. Malekzadeh, R.G. Clegg, A. Cavallaro, H. Haddadi, Protecting sensory data against sensitive inferences, in: *Proc. Workshop on Privacy by Design in Distributed Systems*, 2018.
- [11] G. Vavoulas, C. Chatzaki, T. Malliotakis, M. Pediaditis, M. Tsiknakis, The mobiact dataset: recognition of activities of daily living using smartphones, in: *Proc. International Conference on Information and Communication Technologies for Ageing Well and e-Health*, 2016.
- [12] T.T. Ngo, M.A.R. Ahad, A.D. Antar, M. Ahmed, D. Muramatsu, Y. Makihara, Y. Yagi, S. Inoue, T. Hossain, Y. Hattori, OU-ISIR wearable sensor-based gait challenge: age and gender, in: *Proc. International Conference on Biometrics*, 2019.
- [13] S. Liu, W. Shao, T. Li, W. Xu, L. Song, Recent advances in biometrics-based user authentication for wearable devices: a contemporary survey, *Digit. Signal Process.* (2021) 103120.
- [14] J. Mantyjarvi, M. Lindholm, E. Vildjiounaite, S.-M. Makela, H. Ailisto, Identifying users of portable devices from gait pattern with accelerometers, in: *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2005.
- [15] S. Sprager, M.B. Juric, Inertial sensor-based gait recognition: a review, *Sensors* 15 (9) (2015) 1–39.
- [16] M. Gadaleta, M. Rossi, IDNet: smartphone-based gait recognition with convolutional neural networks, *Pattern Recognit.* 74 (2018) 25–37.
- [17] R. Tolosana, R. Vera-Rodriguez, J. Fierrez, J. Ortega-Garcia, DeepSign: deep on-line signature verification, *IEEE Trans. Biom., Behav., Identity Sci.* 3 (2) (2021) 229–239.
- [18] R. Tolosana, P. Delgado-Santos, A. Perez-Urbe, R. Vera-Rodriguez, J. Fierrez, A. Morales, DeepWriteSYN: on-line handwriting synthesis via deep short-term representations, in: *Proc. AAAI Conference on Artificial Intelligence*, 2021.
- [19] J.M. Ackerson, R. Dave, J. Seliya, Applications of recurrent neural network for biometric authentication & anomaly detection, *Information* 12 (7) (2021) 1–20.
- [20] Q. Zou, Y. Wang, Q. Wang, Y. Zhao, Q. Li, Deep learning-based gait recognition using smartphones in the wild, *IEEE Trans. Inf. Forensics Secur.* 15 (2020) 3197–3212.
- [21] Y. Iwasawa, K. Nakayama, I. Yairi, Y. Matsuo, Privacy issues regarding the application of DNNs to activity-recognition using wearables and its countermeasures by use of adversarial training, in: *Proc. International Joint Conference on Artificial Intelligence*, 2017.
- [22] D. Zhang, L. Yao, K. Chen, Z. Yang, X. Gao, Y. Liu, Preventing sensitive information leakage from mobile sensor signals via integrative transformation, *IEEE Trans. Mob. Comput.* (2021) 1–1.
- [23] G. Garofalo, D. Preuveneers, W. Joosen, Data privatizer for biometric applications and online identity management, *Privacy Identity Manag.* (2019) 209–225.
- [24] P. Terhörst, N. Damer, F. Kirchbuchner, A. Kuijper, Unsupervised privacy-enhancement of face representations using similarity-sensitive noise transformations, *Appl. Intell.* 49 (8) (2019) 3043–3060.
- [25] P. Terhörst, N. Damer, F. Kirchbuchner, A. Kuijper, Suppressing gender and age in face templates using incremental variable elimination, in: *Proc. International Conference on Biometrics*, 2019.
- [26] A. Morales, J. Fierrez, R. Vera-Rodriguez, R. Tolosana, SensitiveNets: learning agnostic representations with application to face images, *IEEE Trans. Pattern Anal. Mach. Intell.* 43 (6) (2020) 2158–2164.
- [27] E. Maiorana, EEG-based biometric verification using Siamese CNNs, in: *Proc. International Conference on Image Analysis and Processing*, 2019.
- [28] J. Johnson, A. Alahi, L. Fei-Fei, Perceptual losses for real-time style transfer and super-resolution, in: *Proc. European Conference on Computer Vision*, 2016.
- [29] A. Boutet, C. Frindel, S. Gamba, T. Jourdan, R.C. Ngueveu, DYSAN: dynamically sanitizing motion sensor data against sensitive inferences through adversarial networks, in: *Proc. ACM Asia Conference on Computer and Communications Security*, 2021.
- [30] O. Hajihassnai, O. Ardakanian, H. Khazaei, ObscureNet: learning attribute-invariant latent representation for anonymizing sensor data, in: *Proc. International Conference on Internet-of-Things Design and Implementation*, 2021.