Balancing Activity Recognition and Privacy Preservation with a Multi-objective Evolutionary Algorithm

Angelica Poli[1], Angela M. Muñoz-Antón[2], Susanna Spinsante[1], and Francisco Florez-Revuelta[3](✉)

[1] Dipartimento di Ingegneria dell'Informazione, Università Politecnica delle Marche, 60131 Ancona, Italy
{a.poli,s.spinsante}@staff.univpm.it

[2] University of Alicante, 03690 Alicante, Spain
mma134@gcloud.ua.es

[3] Department of Computing Technology, University of Alicante, 03690 Alicante, Spain
francisco.florez@ua.es

Abstract. With the widespread of miniaturized inertial sensors embedded in wearable devices, an increasing number of individuals monitor their daily life activities through consumer electronic products. However, long-lasting data collection (e.g., from accelerometer) may expose the users to privacy violations, such as the leakage of personal details. To help mitigate these aspects, we propose an approach to conceal subject's personal attributes (i.e., gender) while maximizing the accuracy on both the monitoring and recognition of human activity. In particular, a Multi-Objective Evolutionary Algorithm (MOEA), namely the Non-dominated Sorting Genetic Algorithm II (NSGA-II), is applied to properly weight input features extracted from the raw accelerometer data acquired with a wrist-worn device (Empatica E4). Experiments were conducted on a large-scale and real life dataset, and validated by adopting the Random Forest algorithm with 10-fold cross validation. Findings demonstrate that the proposed method can highly limit gender recognition (from 89.37% using all the features to 64.38% after applying the MOEA algorithm) while only reducing the accuracy of activity recognition by 5.45% points (from 89.59% to 84.14%).

Keywords: Human Activity Recognition · Privacy preservation · Gender user identification · Wearable sensors · Multi-objective evolutionary algorithm

The support of the *More Years Better Lives JPI*, the Italian Ministero dell'Istruzione, Università e Ricerca (CUP: I36G17000380001) and the Spanish Agencia Estatal de Investigación (grant no: PCIN-2017-114), for this research activity carried out within the project PAAL - Privacy-Aware and Acceptable Lifelogging services for older and frail people, (JPI MYBL award number: PAAL_JTC2017) is gratefully acknowledged.

I. M. Pires et al. (Eds.): GOODTECHS 2021, LNICST 401, pp. 3–17, 2021.
https://doi.org/10.1007/978-3-030-91421-9_1

1 Introduction

Different sensors and sensing modalities can be adopted in *lifelogging* applications, which may expose the subject to risks associated with privacy violations, because of the pervasive and long-lasting collection of data [9,23,25]. Additionally, the potential threat to privacy is also determined by the subject's awareness about the undergoing collection of sensor data, while performing daily life activities. In fact, while the explicit user permission is typically requested, e.g. by mobile devices operating systems, to access and activate data collection from sensors such as cameras, microphones and positioning systems, which are generally perceived as information rich and intrusive of the user's privacy, other sensors such as accelerometers, gyroscopes and barometers are usually considered less dangerous in terms of privacy implications, if not taken into consideration at all [24]. Unfortunately, the vast majority of the accelerometers embedded in consumer wearables and smartphones are usually accessible to several third parties (manufacturers, service providers, developers of apps running on the hosting device), that are to be wisely considered as potentially untrusted if not malicious, especially because the user does not have any control on their capabilities to access the sensor data. As acceleration data collected from wearable devices is shown to enable the leakage of personal details [20,42], suitable approaches to privacy by design and by default in wearable devices, which are prescribed by privacy regulations such as the European GDPR [15], are needed [17].

1.1 Activity Recognition with Accelerometry

Human Activity Recognition (HAR) targets to automatically detect and predict human behaviors through motion data collection. Over the past years, both video- and sensors-based HAR have been explored in many application fields, such as ambient assisted living (AAL) [38], anomaly living patterns [16], fall detection [31], sport [47] and pandemic emergency like COVID-19 [19].

Generally, the development of a HAR system follows a specific sequence of steps, named Activity Recognition Chain (ARC), leading from collected raw data to activity recognition system performance. As detailed in [5], ARC consists of four consecutive steps that include the pre-processing, the segmentation, the feature extraction, and finally the selection of features and classifier training.

Within the context of sensor-based HAR, thanks to the quick spread of miniaturized inertial sensors (Inertial Measurement Units—IMUs) embedded in wearable electronic devices, motion-related raw data can be easily acquired from several consumer electronics products (e.g. smartwatches and smartphones), although their metrological properties are often unavailable [10]. Acceleration data are widely employed for HAR analysis [1,40]: such approaches offer an adequate, comfortable and affordable solution to directly measure the movement of the human body. Several studies focused on monitoring human activities have been reviewed in [13] and [39], where the accelerometer is used either alone or combined with other inertial sensors like gyroscopes and magnetometers.

Acceleration samples generated by the sensor capturing the current user's status may be continuously gathered from the wearable devices, processed and then used to classify the performed activities through the use of Machine Learning (ML) and Artificial Intelligence (AI) algorithms. Among others, some of the most commonly implemented classification algorithms include Naive Bayes (NB), Random Forest (RF), Decision Tree (DT), Support Vector Machine (SVM), k-Nearest Neighbor (kNN), Hidden Markov Model (HMM), Artificial Neural Network (ANN), Convolutional Neural Network (CNN) etc. [22,37]. However, different algorithms can be suitable for different scenarios of HAR, depending on the nature of the data collected [36] and also on the target item to classify (activity, gesture, pose, and so on). All of the mentioned algorithms exploit features computed on the acceleration data, either in the time or in the frequency domain; however, some of the features may capture not only the information strictly needed to classify the activity, but also personal details about the subject performing it. As a consequence, the risk is that the personal information is exposed with the HAR classification results.

1.2 Privacy Preservation

The very rapid growth in the adoption and popularity of wearable devices has nurtured the development of a huge amount of different applications exploiting not only the data collected by means of the on board sensors, but also the personal information each user may decide to provide, such as age, height, gender and weight.

For the latter, several mechanisms and approaches have been proposed in the literature, in order to guarantee anonymous sharing and avoidance of privacy disclosure [7,29,46]. However, the same personal information could be disclosed also without the explicit consensus of the user, by processing some of the motion-related data collected from wearable sensors, especially those generated by accelerometers. As explained in [43], for applications targeting HAR based on acceleration data, a potential privacy leakage exists about the user's identity, because it is possible to infer gait characteristics depending on a user's muscle growth, bone structure, height, and weight, from the sensed accelerometer data. Also, the authors show how feature selection and sampling rate adjustment have an impact on the accuracy of both activity and identity recognition, and they should be considered on the basis of their *mutual information* metric. Based on the above observations, motion data can be classified as a *quasi-identifier* [24], because it may allow the identification and tracking of a user. In fact, every individual has a distinctive way of walking, which is the reason why gait can be a key element of biometric techniques to authenticate and/or identify the user of a wearable device. Moreover, motion data can be drawn from wearables (such as smartbands) and mobile devices without the user's conscious participation or explicit permission, thus becoming a potential threat despite being generally thought to be harmless. Boutet et al. in their report [4] present a privacy-preserving tool to sanitize motion sensor data against unwanted sensitive inferences (thus improving privacy), while keeping an acceptable accuracy of the HAR (thus maintaining data utility). To do so, the tool builds several models to sanitize the motion data against the specified attribute (such as gender), by exploiting Generative Adversarial Networks (GANs). The generated sanitized data may result to be distorted with regards to the raw one, for the aim of the target application (HAR) so, based on a utility/privacy trade-off, several models may be necessary before finding the one that optimizes HAR performance loss, sensitive inference reduction, and data distortion. Authors present test results on available data collections, for which gender inference is reduced up to 41% while decreasing the HAR accuracy only by 3%.

While gait has been proved to enable a subject's identification from his/her motion data in the literature, in a previous paper we discussed the risk of exposure of personal information also when the collected acceleration data are not related to gait, but associated to different types of activities or gestures [34]. Taking all this into consideration, the main objective of this paper is to design a method to properly select features computed from the acceleration data acquired with a smart wristband while performing different activities, and used by supervised classification algorithms, so that HAR accuracy (i.e. the utility of the data) is not affected, while personal attributes (i.e. unnecessary private details in the processed data), such as gender, are concealed. A Multi-Objective Evolutionary Algorithm (MOEA) is applied to find appropriate weights for each feature, differently from the paper cited above which exploits GANs. The remainder of this paper is organized as follows. In Sect. 2 the initial method applied for HAR and gender recognition is presented. This Section also provides a description of the collected dataset, highlighting the features extracted from the accelerometer signal, which serve as input for the recognizing systems. In Sect. 3 a multi-objective evolutionary algorithm is described to find the relevance of each feature so that HAR accuracy percentage is maintained (or improved), while gender recognition is worsened to preserve the private information. Finally, Sect. 5 discusses the approach proposed in the paper and presents some future works.

2 Activity and Gender Recognition

2.1 The Dataset

This work employs the publicly available PAAL ADL Accelerometry dataset [32], a dataset acquired with a wearable multi-sensor device, the Empatica E4 [14], which provides raw activity motion data in real-life conditions.

To the aim of this work, among the signals collected by the sensors embedded in Empatica E4, only the acceleration has been extracted to monitor the users performing different activities of daily living. To promote the real-life acquisition procedure, subjects acted in their natural environment, with no instructions about how and for how long to perform each activity.

The dataset includes 24 different activities performed using real objects. Each activity was repeated between 3 and 5 times by 33 healthy subjects, characterized by a gender balance (19 females and 14 males), and a large age range (between 18 and 77 years, mean = 45.24 years and standard deviation = 18.24 years).

The list of activities is presented in Table 1. This dataset also includes information about the gender and age of each subject.

Table 1. List of activities performed by the subjects.

1. Drink water	7. Take off jacket	13. Sit down	19. Sneeze/cough
2. Eat meal	8. Put on jacket	14. Stand up	20. Blow nose
3. Open a bottle	9. Put on a shoe	15. Writing	21. Washing hands
4. Open a box	10. Take off a shoe	16. Phone call	22. Dusting
5. Brush teeth	11. Put on glasses	17. Type on a keyboard	23. Ironing
6. Brush hair	12. Take off glasses	18. Salute	24. Washing dishes

2.2 Signal Processing and Features Extraction

Each time series of acceleration samples A was filtered by applying a 4^{th} order low-pass Butterworth filter (cut-off frequency: 15 Hz) for eliminating the high frequency noise and preserving the human activities and gestures [35]. In addition, a 3^{rd} order median filter was applied to remove abnormal spikes. In order to infer the information contained in the human activity data, specific features were computed in both time- and frequency-domain, and then extracted, either from each spatial direction (A_x, A_y, A_z) and from the signal magnitude vector (SMV) of acceleration (defined as $SMV_i = \sqrt{A_{x,i}^2 + A_{y,i}^2 + A_{z,i}^2}$, where i is the index of the sample).

A successful technique for extracting features from accelerometer data is to divide the signals into windowed segments [28]. Therefore, we extract features from both raw data of the three axes and SMV, each segmented by fixed-size sliding windows of 5 s (i.e. 160 samples), with 20% (i.e. 1 s) of overlapping between two adjacent windows. Concerning the time domain analysis, features strictly related to the changes in the acceleration signal were evaluated. The frequency domain analysis was conducted by computing the magnitude of discrete Fast Fourier Transform (FFT) on the acceleration signals. The selected features were then extracted from the obtained signals. The list of 62 features used is presented in Table 2 [2,33].

2.3 Implementation

The overall implementation was made using the Python library Scikit-learn. Among the ML classifiers commonly selected to train and test the classification model, Random Forest (RF) was used thanks to its good performance reported in previous similar studies [18,44]. RF is designed to obtain accurate predictions by constructing multiple trees (named also estimators), where each tree is grown with a randomized subset of features. In this study, the whole set of features was used to build each tree of RF, measuring the quality of split with the information gain function. The number of trees was investigated by selecting different

Table 2. List of features extracted in time and frequency domain.

Domain	Features	Computation
Time	Mean	X, Y, Z axes, SMV
	Median	X, Y, Z axes, SMV
	Standard Deviation	X, Y, Z axes, SMV
	Maximum	X, Y, Z axes, SMV
	Minimum	X, Y, Z axes, SMV
	Range	X, Y, Z axes, SMV
	Axes Correlation	XY, YZ, ZX axes
	Signal Magnitude Area	SMV
	Coefficient of Variation	X, Y, Z axes, SMV
	Median Absolute Deviation	X, Y, Z axes, SMV
	Skewness	SMV
	Kurtosis	SMV
	Autocorrelation	SMV
	Percentiles (20^{th} - 50^{th} - 80^{th} - 90^{th})	SMV
	Interquartile range	SMV
	No. of Peaks	SMV
	Peak - Peak Amplitude	SMV
	Energy	SMV
	Root Mean Square	SMV
Frequency	Spectral Entropy	SMV
	Spectral Energy	SMV
	Spectral Centroid	SMV
	Mean	X, Y, Z axes, SMV
	Standard Deviation	X, Y, Z axes, SMV
	Percentiles (25^{th} - 50^{th} - 75^{th})	SMV

numbers of estimators (from 10 to 170) to verify whether such parameter can improve the model accuracy for activity and gender recognition.

In order to train the model, the k-fold cross validation was implemented with $k = 10$. This means that the original dataset was randomly partitioned into k subsets: $k-1$ as training set, and the remaining as validation set. After k times, the average of the k performance measurements on the k validation sets gives the cross-validated performance.

2.4 Results

Generally, in the HAR context, supervised learning approaches are widely exploited by using labeled data as inputs for predicting the classification of unknown data through the ML algorithms. In this section, we present the results of the RF supervised approach used to investigate both the activity and gender recognition, separately.

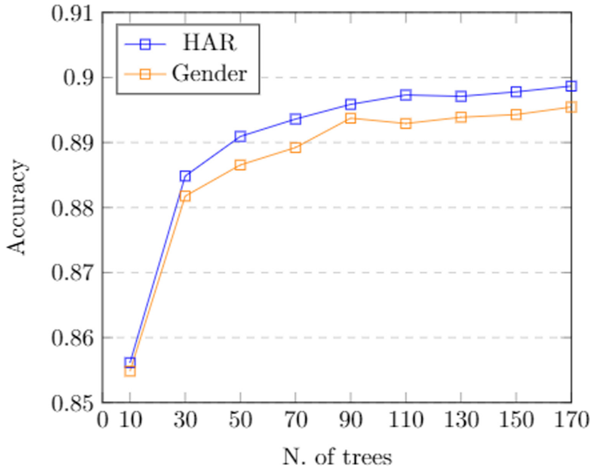

Fig. 1. Number of trees for RF classifier.

The experimental results are shown in Fig. 1, where the number of trees is associated to a value of accuracy percentage achieved by the RF algorithm. Contrary to what often may be thought, the results show that a large number of trees does not necessarily cause a better recognition performance for the RF algorithm. In fact, for both models, the percentages of improvement stabilize or slightly decrease as the number of trees increases, especially after a number of trees equal to 90. This means that the benefit in recognition performance from using more trees is lower than the cost in computation time for learning the additional trees. Therefore, at 90 estimators, the global accuracy percentage is 89.59% and 89.37% for human activity and gender recognition, respectively. Besides the accuracy, the performance of the activity and gender recognition was assessed using the confusion matrix, that summarizes the information about the actual and predicted classes, either correctly or wrongly classified. Figure 2 shows the results obtained for the RF algorithm to classify each activity. As it is clear, the highly misclassified activities are the following ones: open a box classified as open bottle, take off a shoe classified as put on a shoe, take off glasses classified as put on glasses, and stand up classified as sit down. On the other hand, Fig. 3 shows the performance obtained for the RF algorithm in distinguishing the gender of the participants. Male subjects were confused more than female ones, probably because there are more female users in the dataset.

3 Concealing Private Information with a Multi-objective Evolutionary Algorithm

The above results can be considered good for HAR. However, if the goal is to protect identity details of the users performing those activities, this set of features also gets good results for gender recognition. Then, the question is:

Actual class \ Predicted class	drink water	eat meal	open bottle	open a box	brush teeth	brush hair	take off jacket	put on jacket	put on a shoe	take off a shoe	put on glasses	take off glasses	sit down	stand up	writing	phone call	type on a keyboard	salute	sneeze/cough	blow nose	washing hands	dusting	ironing	washing dishes
drink water	90	1									3	3	1				1		1	2				
eat meal	2	55	5	4							5	3	6	7		3			5	5				
open bottle	1	4	44	30					1	1	1	3	3	1		1		1	1	9				
open a box	1	1	24	57						1	2	1		1				1	1	9				
brush teeth					95	0	0	0	0	1							1				0	0	2	0
brush hair	0				1	90	1	2	2	0				0				0	0		0	1	2	0
take off jacket						1	84	9	1	1								2				1		
put on jacket						1	3	87	4	0					0							2	1	1
put on a shoe			0		0	1	0	0	85	3			0				0			1	1	1	4	3
take off a shoe		1	3	1	1	0	1		14	51	0		4	2	0	0	0	1	3	3	2	3	4	3
put on glasses	5	3		1							70	11	1	3		6		1	1					
take off glasses	5	3	3	6						1	12	54	4	2		1		4	3	1				
sit down		0	2	1						1	0	1	72	16				1	1	1				1
stand up		4	2	1						2		1	18	69				1	1	1				
writing															98		1						0	
phone call		3		2		1	2			1	10	3	1	2		73		1		1				1
type on a keyboard									0						3		96			0			1	
salute				1		1	1	1		1	1	3				3		90						
sneeze/cough	3	6	6	8						2	8	5	5	6		1		6	41	3		1	1	
blow nose	3	4	4	5	1	1			1	1	1	1	2		1	1	1		4	67		1	1	2
washing hands		0					0	0	1	1			1	0							92	0	0	4
dusting				0	0	1	1	1	3	2						0		1		0	1	85	2	2
ironing					0	0			0	0	0		0		0		2		0	0		0	96	1
washing dishes					0	0	0		1	0										0	2	1	1	95

Fig. 2. Confusion matrix for HAR recognition.

Actual \ Predicted	male	female
male	79	21
female	4	96

Fig. 3. Confusion matrix for gender recognition.

Can the input features be transformed so that HAR remains good but gender (or any other private data) recognition accuracy decreases? This would lead to systems that perform well for HAR, maintaining the utility of the data, but, simultaneously, conceal private information. This work presents a mechanism to filter the input features in order to obtain this objective. We propose the use of a multi-objective evolutionary algorithm to find appropriate weights for each feature.

3.1 Evolutionary Feature Weighting

This is a similar objective to previous works on evolutionary feature subset selection or feature weighting. This is, trying to approximate the optimal degree of influence of individual features using a training set. However, the final goal of this work is to filter the signal captured by the wristband trying to maximize the

HAR accuracy while at the same time to minimize the recognition of personal characteristics of the subjects, e.g. gender or age.

In the case of evolutionary feature weighting, the approach is to consider an individual as a real vector where each gene represents the weight given to each feature. Two main models are presented in the literature to implement this [6, 45]: the filter model and the wrapper model. The former selects the features based on a priori decisions on feature relevance according to some measures (information, distance, dependence or consistency) [30], ignoring the learning algorithm underneath. In the latter, the feature selection algorithm exists as a wrapper around the learning algorithm; in the search of a feature subset, the learning algorithm itself is used as part of the evaluation function [21]. The main disadvantage of the wrapper approach is the time needed for the evaluation of each feature weight [26]. On the other hand, filter-based approaches, although faster, find worse solutions in general [6].

This paper follows the wrapper approach, trying to find appropriate weights for the features presented in Table 2 using RF as classifier.

3.2 Multi-objective Evolutionary Algorithms

In multi-objective optimization, the goal is to optimize simultaneously several objective functions. These different functions have conflicting objectives, i.e., optimizing one affects the others. Therefore, there is not a unique solution but a set of solutions. The set of solutions in which the different objective components cannot be simultaneously improved constitute a Pareto front. Each solution in the Pareto front represents a trade-off between the objectives. MOEAs [8] are heuristic algorithms to solve problems with multiple objective functions. The three goals of an MOEA are [41]: (1) to find a set of solutions as close as possible to the Pareto front (convergence); (2) to maintain a diverse population that contains dissimilar individuals to promote exploration and to avoid poor performance due to premature convergence (diversity); and (3) to obtain a set of solutions that spreads in a more uniform way over the Pareto front (coverage).

This work employs the Non-dominated Sorting Genetic Algorithm II (NSGA-II) [11] as wrapper algorithm. NSGA-II has the three following features: (1) it uses an elitist principle, i.e., the elites of a population are given the opportunity to be carried to the next generation; (2) it uses an explicit diversity preserving mechanism (Crowding distance); and (3) it emphasizes the non-dominated solutions. The algorithm will obtain a set of solutions, some of them optimizing one over the other objective and vice versa. From these set of solutions, a specific solution fulfilling particular conditions could be selected.

3.3 Characteristics of the Algorithm

An individual in the population (potential solution) is encoded as a real vector U whose elements $u_j, \forall j \in [1..62]$ represent the weight of a particular feature during the classification (62 being the number of features, see Table 2), i.e. each feature is multiplied by the appropriate weight before being input to the classifier.

The fitness functions to be optimized are:

$$f_1 = HAR \tag{1}$$

$$f_2 = \left| Gender\ Recognition - \frac{1}{2} \right| \tag{2}$$

The objective will be to maximize f_1 while minimizing f_2, i.e. maximizing HAR while taking gender recognition close to random. Therefore, for other private information, the second term in f_2 must be $\frac{1}{Number\ of\ categories}$.

4 Results

This work has employed the implementation of NSGA-II offered by pymoo [3], a multi-objective optimization framework in Python, using the following parameters, which have been selected experimentally:

- Size of the population: 50;
- New individuals (offsprings) created per generation: 10; and
- Number of generations without changes in the best individual to stop the algorithm: 100

Some other characteristics of the algorithm are:

- The individuals in the initial population are created with random real values between 0 and 1;
- Binary tournament is used to select the parents to generate a new offspring;
- Simulated Binary Crossover (SBX) [12] with default parameters is employed to create each individual;
- Each new individual is mutated by applying Polynomial Mutation [12] with default parameters; and
- Duplicates are eliminated after merging the parent and the offspring population. If there are duplicates with respect to the current population or in the offsprings itself they are removed and the mating process is repeated to fill up the offsprings until the desired number of unique offsprings is met.

The application of this algorithm obtains a set of solutions, some of them optimizing one over the other objective and vice versa (Fig. 4). From these set of solutions, this work selects the solution closer (using Euclidean distance) to perfect HAR (value equal to 1) and random gender recognition (value equal to 0.5). The accuracy results of this selected final solution (marked in green in the Figure) are 84.14% and 64.38%, for activity and gender recognition, respectively. Comparing with the results presented in Sect. 2.4 (in red in the Figure), in which all the features are equally considered, HAR accuracy has been reduced 5.45% points while gender recognition has worsened 24.99.

Figure 5 shows that results are similar to those shown in Fig. 2, although accuracy has been reduced for most of the classes. However, gender classification has been dramatically affected by applying the MOEA as almost all the inputs are now identified as being performed by a female user (Fig. 6).

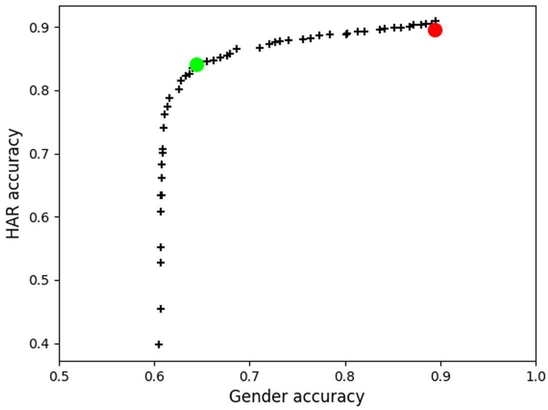

Fig. 4. Final set of solutions for a run of the MOEA algorithm. The best balanced solution is marked in green. The initial result in which all features are equally considered (presented in Sect. 2.4) is marked in red. For the sake of clarity, 0.5 has been added to f_2, which is then the gender accuracy. (Color figure online)

	Predicted class																							
Actual class	drink water	eat meal	open bottle	open a box	brush teeth	brush hair	take off jacket	put on jacket	put on a shoe	take off a shoe	put on glasses	take off glasses	sit down	stand up	writing	phone call	type on a keyboard	salute	sneeze/cough	blow nose	washing hands	dusting	ironing	washing dishes
drink water	83	1	1								7	4	1	1	1		1			1			1	
eat meal	3	43	7	8					1	1	3	5	7	7	1	3		1	5	4			2	
open bottle	1	5	41	26					2	2	3	1	3	1	2		1	1	3	5		1		
open a box	1	1	27	50					1		2	1	3	3	1	1	1	1	2	3		1		
brush teeth					91			0	0	0					1		2				0	0	4	1
brush hair	0				3	79	0	2	2	0	0	0				0		0		0	0	1	7	3
take off jacket					0	1	74	18	1	0								1			0	1	0	1
put on jacket						1	1	84	2						1		0	0				3	1	5
put on a shoe					2	1	0	0	68	1			1	0	1		1				2	1	6	16
take off a shoe		1	5	0	7	1	1		9	16	0		3	3	9	1	1	2	2	0	5	1	11	21
put on glasses	6	3	1	1							63	12	2			8		3	1				1	
take off glasses	9	2	5	6							14	42	6	2	1	2	1	7	2				1	
sit down		2	2	3		1			1	2	1	2	65	16	1	2		1	2			1	1	
stand up		5	1	1						1	1	1	13	70	1	2		1	2			1	1	
writing															98		2						0	
phone call	1	3		1		1	2		1		7	3	2	1		67		5		1	1	4	1	1
type on a keyboard															6		93						1	
salute		1		1		1	1	1			1	4	1			6		82				1		
sneeze/cough	2	3	5	5		1				1	10	5	5	6	3	1	1	8	36	3		1	3	
blow nose	1	2	6	4	3	1				1	1	1	2	1	6	1	6		1	48		1	12	5
washing hands		0			0			0	0	0			0	0	0					0	84	0	0	13
dusting					3	1	2	3	0	1						0		0			1	72	5	13
ironing					1				0	0					1	0	3		0	0	0	0	92	1
washing dishes					1			0	0								0				1	0	1	96

Fig. 5. Confusion matrix related to the HAR.

		Predicted	
		male	female
Actual	male	10	90
	female	0	100

Fig. 6. Confusion matrix related to the gender.

5 Conclusion

In this work, the authors propose a method for privacy preservation in non gait-related human motion data collection. More precisely, an MOEA is applied on a dataset of 24 different activities of daily living performed in realistic conditions by 33 healthy subjects, and for which acceleration signals were collected from the medical grade wrist worn device Empatica E4.

In order to investigate how much private information (i.e. related to gender) is included in the set of features extracted from accelerometer data, the RF algorithm was firstly adopted for both the activity and gender recognition separately. Then, the NSGA-II was employed as a wrapper algorithm to find the appropriate weights for each feature so that HAR accuracy is maximized while simultaneously gender accuracy is minimized. In order to validate the MOEA with respect to the initial method, the classification accuracy along with the confusion matrices were investigated.

Initial results show a global accuracy percentage of 89.59% and 89.37% for human activity and gender recognition, respectively, by adopting the RF algorithm with a number of trees equal to 90. In particular, the confusion matrix for HAR suggests that some activities are highly misclassified because they are difficult to distinguish using only a single accelerometer worn on the wrist, as the movement of the device may be similar in some contrary activities (e.g. stand up and sit down). Also, the same activity performed with two different tools was harder to distinguish (e.g. open a box and open bottle). Regarding gender recognition, the results demonstrate that the user's gender can be easily detected from the wrist motion data. Finding the appropriate weights by using an MOEA allows to conceal the gender of the user while HAR is not considerably affected, resulting in a global accuracy of 84.14% and 64.38% for human activity and gender recognition, respectively. Additionally, Fig. 4 shows that the output of the multi-objective evolutionary algorithm is not a single set of weights (single solution) but a set of solutions constituting a Pareto front. Therefore, under some conditions instead of choosing the best balanced solution (as this work proposes) any other of the obtained solutions could be chosen, for instance in the case that, for a specific user, activity recognition needs to be prioritized over privacy protection. Getting back to the motivations of this study, implementing a concealing strategy on-board the sensor based on the selected condition on the Pareto front, could be seen as a way to implement what is envisaged as a privacy-by-design approach by the regulations.

In the future, this work will be extended to consider other private information which users would like to conceal, e.g. age. In that case, in which more than two objective functions were to be considered, the optimization will be carried out with a many-objective evolutionary algorithms [27]. Additionally, other set of features extracted from the accelerometer data will be explored. For instance, extracting deep features from the original data instead of using the handcrafted features presented in Table 2.

References

1. Balli, S., Sağbaş, E.A., Peker, M.: Human activity recognition from smart watch sensor data using a hybrid of principal component analysis and random forest algorithm. Meas. Control **52**(1–2), 37–45 (2019)
2. Bautista-Salinas, D., González, J., Méndez, I., Mozos, O.: Monitoring and prediction of mood in elderly people during daily life activities. In: 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference (EMBC), vol. 2019, pp. 6930–6934. IEEE (2019)
3. Blank, J., Deb, K.: Pymoo: multi-objective optimization in python. IEEE Access **8**, 89497–89509 (2020)
4. Boutet, A., Frindel, C., Gambs, S., Jourdan, T., Ngueveu, R.C.: DYSAN: dynamically sanitizing motion sensor data against sensitive inferences through adversarial networks. In: INRIA Informatics Mathematics, pp. 1–26. INRIA (2020)
5. Bulling, A., Blanke, U., Schiele, B.: A tutorial on human activity recognition using body-worn inertial sensors. ACM Comput. Surv. (CSUR) **46**(3), 1–33 (2014)
6. Cantú-Paz, E.: Feature subset selection, class separability, and genetic algorithms. In: Deb, K. (ed.) GECCO 2004. LNCS, vol. 3102, pp. 959–970. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24854-5_96
7. Chen, C.M., Xiang, B., Wu, T.Y., Wang, K.H.: An anonymous mutual authenticated key agreement scheme for wearable sensors in wireless body area networks. Appl. Sci. (Switzerland) **8**(7) (2018). https://doi.org/10.3390/app8071074
8. Coello, C.A.C., Lamont, G.B., Van Veldhuizen, D.A., et al.: Evolutionary Algorithms for Solving Multi-objective Problems, vol. 5. Springer, Heidelberg (2007). https://doi.org/10.1007/978-0-387-36797-2
9. Connor, P., Ross, A.: Biometric recognition by gait: a survey of modalities and features. Comput. Vis. Image Underst. **167**, 1–27 (2018)
10. Cosoli, G., Spinsante, S., Scalise, L.: Wrist-worn and chest-strap wearable devices: systematic review on accuracy and metrological characteristics. Measurement **159**, 107789 (2020)
11. Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.: A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Trans. Evol. Comput. **6**(2), 182–197 (2002)
12. Deb, K., Sindhya, K., Okabe, T.: Self-adaptive simulated binary crossover for real-parameter optimization. In: Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation, GECCO 2007, pp. 1187–1194. Association for Computing Machinery, New York (2007)
13. Elbasiony, R., Gomaa, W.: A survey on human activity recognition based on temporal signals of portable inertial sensors. In: Hassanien, A.E., Azar, A.T., Gaber, T., Bhatnagar, R., F. Tolba, M. (eds.) AMLTA 2019. AISC, vol. 921, pp. 734–745. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-14118-9_72

14. Empatica Inc.: Empatica E4. https://www.empatica.com/en-eu/research/e4/. Accessed 29 Apr 2021
15. European Commission: Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (2016). https://eur-lex.europa.eu/. Accessed 26 July 2021
16. Fahad, L.G., Tahir, S.F.: Activity recognition and anomaly detection in smart homes. Neurocomputing **423**, 362–372 (2021)
17. Flórez-Revuelta, F., Mihailidis, A., Ziefle, M., Colonna, L., Spinsante, S.: Privacy-aware and acceptable lifelogging services for older and frail people: the PAAL project. In: 2018 IEEE 8th International Conference on Consumer Electronics - Berlin (ICCE-Berlin), pp. 1–4, September 2018
18. Gomaa, W.: Comparative analysis of different approaches to human activity recognition based on accelerometer signals. In: Hassanien, A.E., Darwish, A. (eds.) Machine Learning and Big Data Analytics Paradigms: Analysis, Applications and Challenges. SBD, vol. 77, pp. 303–322. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-59338-4_16
19. Hoang, M.L., Carratù, M., Paciello, V., Pietrosanto, A.: Body temperature—indoor condition monitor and activity recognition by MEMS accelerometer based on IoT-alert system for people in quarantine due to COVID-19. Sensors **21**(7) (2021)
20. Jain, A., Kanhangad, V.: Gender classification in smartphones using gait information. Expert Syst. Appl. **93**, 257–266 (2018)
21. John, G.H., Kohavi, R., Pfleger, K.: Irrelevant features and the subset selection problem. In: Machine Learning Proceedings 1994, pp. 121–129. Elsevier (1994)
22. Kilic, S., Kaya, Y., Askerbeyli, I.: A new approach for human recognition through wearable sensor signals. Arab. J. Sci. Eng. **46**, 4175–4189 (2021)
23. Kim, J.W., Lim, J.H., Moon, S.M., Jang, B.: Collecting health lifelog data from smartwatch users in a privacy-preserving manner. IEEE Trans. Consum. Electron. **65**(3), 369–378 (2019). https://doi.org/10.1109/TCE.2019.2924466
24. Kröger, J.: Unexpected inferences from sensor data: a hidden privacy threat in the Internet of Things. In: Strous, L., Cerf, V.G. (eds.) IFIPIoT 2018. IAICT, vol. 548, pp. 147–159. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-15651-0_13
25. Ksibi, A., Alluhaidan, A.S.D., Salhi, A., El-Rahman, S.A.: Overview of lifelogging: current challenges and advances. IEEE Access **9**, 62630–62641 (2021). https://doi.org/10.1109/ACCESS.2021.3073469
26. Lanzi, P.L.: Fast feature selection with genetic algorithms: a filter approach. In: Proceedings of 1997 IEEE International Conference on Evolutionary Computation (ICEC 1997), pp. 537–540. IEEE (1997)
27. Li, B., Li, J., Tang, K., Yao, X.: Many-objective evolutionary algorithms: a survey. ACM Comput. Surv. (CSUR) **48**(1), 1–35 (2015)
28. Li, J., Tian, L., Wang, H., An, Y., Wang, K., Yu, L.: Segmentation and recognition of basic and transitional activities for continuous physical human activity. IEEE Access **7**, 42565–42576 (2019)
29. Liu, F., Li, T.: A clustering k-anonymity privacy-preserving method for wearable IoT devices. Secur. Commun. Netw. **2018** (2018). https://doi.org/10.1155/2018/4945152
30. Liu, H., Motoda, H.: Feature Selection for Knowledge Discovery and Data Mining, vol. 454. Springer, Heidelberg (2012)

31. Montanini, L., Del Campo, A., Perla, D., Spinsante, S., Gambi, E.: A footwear-based methodology for fall detection. IEEE Sens. J. **18**(3), 1233–1242 (2017)
32. Muñoz-Anton, A.M., Poli, A., Spinsante, S., Florez-Revuelta, F.: PAAL ADL Accelerometry dataset (2021). https://doi.org/10.5281/zenodo.4750904
33. Poli, A., Scalise, L., Spinsante, S., Strazza, A.: ADLs monitoring by accelerometer-based wearable sensors: effect of measurement device and data uncertainty on classification accuracy. In: 2020 IEEE International Symposium on Medical Measurements and Applications (MeMeA), pp. 1–6. IEEE (2020)
34. Poli, A., Strazza, A., Cecchi, S., Spinsante, S.: Identification issues associated with the use of wearable accelerometers in lifelogging. In: Gao, Q., Zhou, J. (eds.) HCII 2020. LNCS, vol. 12207, pp. 338–351. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-50252-2_26
35. Poli, A., Spinsante, S., Nugent, C., Cleland, I.: Improving the collection and understanding the quality of datasets for the aim of human activity recognition. In: Chen, F., García-Betances, R.I., Chen, L., Cabrera-Umpiérrez, M.F., Nugent, C. (eds.) Smart Assisted Living. CCN, pp. 147–165. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-25590-9_7
36. Rafferty, J., Nugent, C.D., Liu, J., Chen, L.: From activity recognition to intention recognition for assisted living within smart homes. IEEE Trans. Hum.-Mach. Syst. **47**(3), 368–379 (2017)
37. Joy Rakesh, Y., Kavitha, R., Julian, J.: Human activity recognition using wearable sensors. In: Satapathy, S.C., Zhang, Y.-D., Bhateja, V., Majhi, R. (eds.) Intelligent Data Engineering and Analytics. AISC, vol. 1177, pp. 527–538. Springer, Singapore (2021). https://doi.org/10.1007/978-981-15-5679-1_51
38. Ranieri, C.M., MacLeod, S., Dragone, M., Vargas, P.A., Romero, R.A.F.: Activity recognition for ambient assisted living with videos, inertial units and ambient sensors. Sensors **21**(3) (2021)
39. Slim, S., Atia, A., Elfattah, M., Mostafa, M.S.M.: Survey on human activity recognition based on acceleration data. Intl. J. Adv. Comput. Sci. Appl **10**, 84–98 (2019)
40. Sousa Lima, W., Souto, E., El-Khatib, K., Jalali, R., Gama, J.: Human activity recognition using inertial sensors in a smartphone: an overview. Sensors **19**(14) (2019)
41. Trivedi, A., Srinivasan, D., Sanyal, K., Ghosh, A.: A survey of multiobjective evolutionary algorithms based on decomposition. IEEE Trans. Evol. Comput. **21**(3), 440–462 (2017)
42. Van hamme, T., Garofalo, G., Argones Rúa, E., Preuveneers, D., Joosen, W.: A systematic comparison of age and gender prediction on IMU sensor-based gait traces. Sensors **19**(13) (2019). https://doi.org/10.3390/s19132945. https://www.mdpi.com/1424-8220/19/13/2945
43. Xiao, F., et al.: An information-aware visualization for privacy-preserving accelerometer data sharing. HCIS **8**(1), 1–28 (2018). https://doi.org/10.1186/s13673-018-0137-6
44. Yacchirema, D.C., de Puga, J.S., Palau, C., Esteve, M.: Fall detection system for elderly people using IoT and ensemble machine learning algorithm. Pers. Ubiquit. Comput. **23**, 801–817 (2019)
45. Yusta, S.C.: Different metaheuristic strategies to solve the feature selection problem. Pattern Recogn. Lett. **30**(5), 525–534 (2009)
46. Zhang, Z., Wang, X., Uden, L., Zhang, P., Zhao, Y.: e-DMDAV: a new privacy preserving algorithm for wearable enterprise information systems. Enterp. Inf. Syst. **12**(4), 492–504 (2018)
47. Zhuang, Z., Xue, Y.: Sport-related human activity detection and recognition using a smartwatch. Sensors **19**(22), 5001 (2019)