Personalized gait detection using a wrist-worn accelerometer

Guglielmo Cola[1], Marco Avvenuti[1], Fabio Musso and Alessio Vecchio[1]

***Abstract*— Wrist-worn devices, such as smartwatches and smart bands, have brought about the unprecedented opportunity to continuously monitor gait during daily routines. However, the use of a single wrist-worn unit for gait analysis is challenging for a variety of reasons. Indeed, the signal collected at the user's wrist is subject to a significant "noise" with respect to other body positions (e.g. waist), mainly due to the arm swing while walking and other unpredictable hand movements. The aim of this paper is to investigate the design and evaluation of a lightweight and reliable gait detection technique for wrist-worn devices. To this end, the proposed method creates a personalized model of the user's gait patterns. The model is created through an automatic training phase, which requires the temporary use of an additional device (smartphone) to gather true gait segments. After, anomaly detection is used to distinguish gait from other activities. Gait data from 20 volunteers have been collected to test and evaluate the proposed technique. Volunteers were asked to walk at different pace, with their normal arm swing or placing the hand inside of a pocket. Results show that the proposed method can reliably distinguish gait from spurious hand movements.**

I. Introduction and Related Work

Wearable devices equipped with inertial sensors are frequently used for monitoring the physical activity of their users. The vast majority of these devices is nowadays represented by smartwatches and wristbands, as they are relatively unobtrusive and easy to wear.

Collecting movements at the user's wrist presents several benefits with respect to previously adopted solutions, which were based on smartphones generally placed at the user's waist (e.g. attached to the belt) or carried in a pocket. First, wearing a device on the wrist is usually considered more acceptable [1]. The unobtrusiveness of these systems actually permits to continuously monitor daily activities and sleep periods without much hassle. The NHANES program carried out a physical activity monitoring study on large scale using wearable devices. Initially, the devices were placed near the hip, but in subsequent years they were moved to the wrist as it resulted in higher wear time [2]. In addition, when the user is at home, the smartphone can be placed elsewhere (in charging, on a table etc.) instead of being carried by the user. There are also some categories of users that tend to bring the smartphone in a bag. A device worn on the wrist is not affected by these types of problems.

Moreover, recognizing some types of activities can be easier when using a wrist-worn device, in particular when these activities make use of hands. The study in [3] presents a personalized method for recognizing activities such as shopping, eating, brushing teeth and typing. Another approach, presented in [4], uses machine learning to identify activities through a wristband. It confirms that it is possible to achieve acceptable performance with a wrist-worn accelerometer. The experiments in [5] demonstrate that, for certain activities performed in domestic environments, the model based on data collected by a wrist-worn device provides better results.

Nevertheless, the use of wrist-worn devices also presents some problems. In particular, due to the lack of a relatively fixed position with respect to the user's center of mass, the recognition of gait can be significantly more challenging than when using other body positions [6].

The aim of this paper is to achieve accurate gait detection with a wrist-worn device embedding an accelerometer. More specifically, we developed a gait detection technique that is accurate (i.e., spurious hand movements do not lead to false recognitions, whereas gait is detected as soon as the user walks a few steps) and lightweight (i.e., it can be executed in real time in miniaturized devices with limited resources, enabling continuous monitoring of gait without hindering battery life). This is achieved through the creation of a personalized gait model obtained through an automatic training phase. The training phase requires the use of a pocket-worn device (e.g. smartphone) in addition to the wrist-worn device, during the first days of use of the system. The pocket-worn accelerometer, thanks to a more favorable position, can help the wrist-worn device in the creation of a training set that includes examples of the user's true gait data. After the training phase, the derived model enables the wrist-worn device to detect gait with high accuracy.

II. Method

The *automatic training phase* of the proposed approach is described in Figure 1. As mentioned above, this phase exploits the availability of a paired device in the user's pocket to create a personalized model of the user's gait patterns. Both the wrist-worn device (WD) and the pocket-worn device (PD) run a *gait segment detection* (GSD) technique. GSD aims to identify a sequence of consecutive steps (gait segment) in the acceleration signal. PD takes advantage of its positioning near the user's center of mass to achieve high accuracy in detecting gait segments, and it can thus pinpoint the starting and ending times of each walking episode (gait intervals) with high accuracy. On the contrary, GSD at wrist position tends to capture spurious hand movements in addition to real gait segments. For this reason, in the proposed approach PD acts as a "supervisor"

*This work was supported by the PRA 2016 project "Analysis of sensory data: from traditional sensors to social sensors", funded by the University of Pisa.

[1]G. Cola, M. Avvenuti and A. Vecchio are with the Dipartimento di Ingegneria dell'Informazione, University of Pisa, 56122 Pisa, Italy g.cola@iet.unipi.it

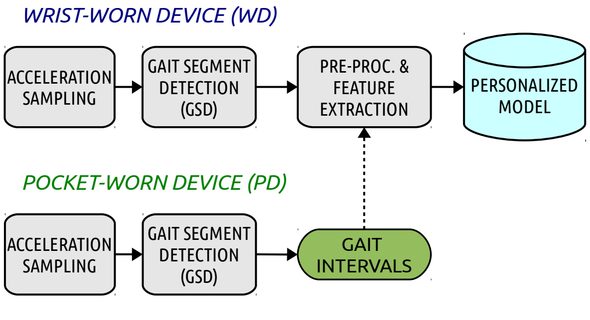

Fig. 1. Automatic training phase

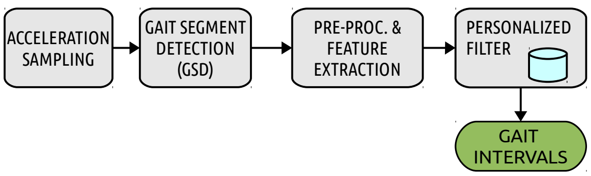

Fig. 2. Independent gait detection at wrist position

during an initial training phase, in order to support WD in the creation of a personalized gait model. In fact, WD exploits PD's gait intervals to label its own gait segments as gait or non-gait – only the former are processed (by the *pre-processing and feature extraction* module) and added to the user's *personalized model*.

In the second phase (Figure 2), after the training procedure has been completed, WD uses the personalized model to autonomously filter out non-gait data and output accurate gait intervals.

In the following we provide more details about GSD, feature-extraction, and the personalized filter.

A. GSD technique

Walking (gait) detection algorithms usually rely on the presence of peaks in the acceleration magnitude signal (for instance [7]–[9]). Such peaks are generated by the foot contact at each step. An example of the acceleration measured during gait at pocket and wrist position is shown in Figure 3. Steps (foot contacts) are highlighted with stripes and numbered. In both traces, the diversity of the signal between consecutive steps is ascribable to the side of the body where the sensor has been worn – the step made with the leg closer to the sensor (red/odd number) is marked by a higher acceleration. Hereafter, we refer to these steps as *dominants*, whereas we refer to the steps producing the lower acceleration as *secondaries*. In the given example, the two sensors were placed on the same side of the user's body, and thus dominants and secondaries in the two signals are aligned. The difference in acceleration amplitude between the two signals is due to the fact that the impact with the ground is transmitted to the sensor through the body, and since the accelerometer placed on the wrist is farther from the foot, the signal is less intense [10]. Also, the presence of arm swing may further reduce the amplitude of wrist acceleration, and almost "hide" some of the secondaries.

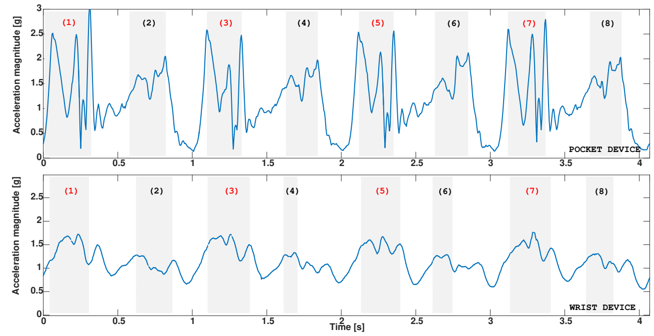

Fig. 3. Gait acceleration collected at the waist and at the wrist

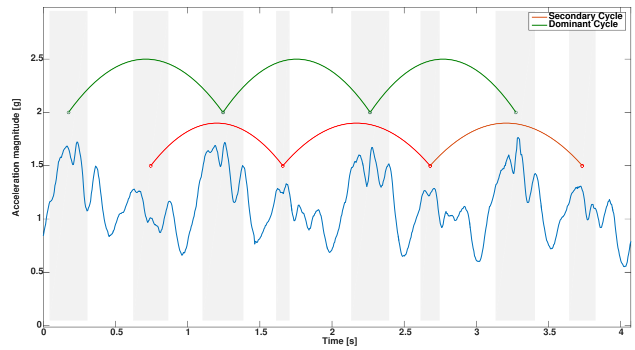

Fig. 4. Gait cycles for dominant and secondary steps

We adapted the algorithm presented in [7] to achieve a high detection rate at wrist position. That algorithm uses a threshold on the acceleration magnitude signal to identify the groups of peaks produced at each step. A gait segment is detected when the following conditions are met: i) a minimum number of consecutive steps are found (e.g. eight steps); ii) a regularity test based on the standard deviation of step duration is passed; iii) an interval of $\sim$1 s without further steps is found, actually terminating the segment.

GSD at wrist was modified as follows: i) the threshold was lowered to ensure the proper detection of dominants and the majority of secondaries; ii) the *gait cycle* detection mode was introduced to manage the above mentioned problem of "hidden" secondaries. More specifically, the technique proposed here is able to automatically switch from step recognition mode to gait cycle mode. A gait cycle is composed of two consecutive steps (Figure 4). Thus, the gait-cycle period is the interval between two consecutive dominants or secondaries. The upper and lower bounds of this interval were found in a reasonable search space, analyzing real gait segments collected from a number of users. Whenever two groups of peaks are detected, the algorithm evaluates if they are separated by a period that falls within the previously found bounds. If so, the gait cycle mode is activated, and since a secondary step might have been missed, the algorithm proceeds by evaluating the regularity of gait cycles instead of the regularity of single steps. In this way, it is possible to detect a gait segment even in the presence of a "hidden" secondary step and preserve the detection rate. On the other

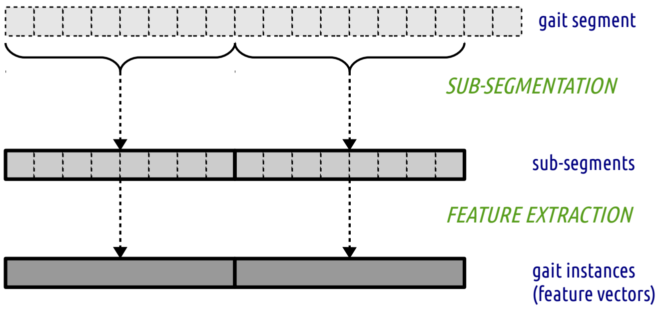

Fig. 5. Data flow diagram, from gait segments to gait instances

hand, the algorithm becomes more likely to detect spurious hand movements, leading to false detections. This specific problem is addressed by the personalized filter.

B. Pre-processing and feature extraction

This module is aimed at preparing the inputs for the personalized filter. The procedure is illustrated by the data flow diagram in Figure 5. A gait segment, identified by means of GSD, is a sequence of eight or more steps. *Sub-segmentation* is used to extract uniform sub-segments with eight steps. In the example in Figure 5 the gait segment is made of 18 steps. Hence, two sub-segments are found, while the last two steps in the segment are ignored. Each sub-segment is then processed with feature extraction algorithms. The result is a distinct feature vector for each sub-segment. Hereafter, we refer to such feature vectors as *gait instances* – a gait instance is the input of the personalized filter, and it is either classified as gait or non-gait. The remaining steps at the end of the original gait segment are classified according to the result obtained by the closest gait instance.

During the training phase, pre-processing and feature extraction are applied only to the sub-segments that are confirmed as real gait by PD. The resulting gait instances are added to the personalized model. Non-gait sub-segments are simply discarded. When training is over, WD is independent and each detected segment is processed to find gait instances. Gait instances, in turn, are classified as gait or non-gait exploiting the personalized filter described in the next subsection.

The rationale behind using sub-segments instead of entire segments as inputs to the personalized filter is twofold: i) by using sub-segments of eight steps, the system is able to classify long segments containing gait and non-gait data with higher accuracy; ii) the use of a predefined and limited number of steps enables the execution of the method in real time on devices with limited resources (e.g. microcontroller class devices with $\sim$ 16 KB of RAM).

The features extracted from each sub-segment are listed in Table I. The list includes well-known statistical features such as mean, median, skewness, root mean square (RMS), and mean crossing rate (MCR). The average absolute acceleration variation (AAV) is the average of the absolute variation of consecutive acceleration samples. It has been previously used in fall detection and gait analysis [11], [12].

TABLE I
LIST OF FEATURES.

AC_C1	AC_DP2	
AAV_y	MCR_m	MCR_z
$mean_x$	$mean_y$	$mean_z$
$median_m$	$median_x$	$median_z$
RMS_z	$skewness_m$	$skewness_y$

Autocorrelation-based features are used to evaluate gait periodicity and regularity. Unbiased autocorrelation coefficients are found as follows:

$$AC_k = \frac{1}{N-k}\sum_{i=1}^{N-k} r_i * r_{i+k} \tag{1}$$

where AC_k is the k-th unbiased autocorrelation coefficient; N is the number of acceleration samples in the sub-segment; r_i is the i-th acceleration magnitude sample minus the mean of the magnitude samples in the sub-segment. The autocorrelation function returns a sequence of autocorrelation coefficients that have peak values in respect of lags equivalent to the periodicity of the signal. These peak values are called *dominant periods*. AC_DP1 and AC_DP2 describe the lag of the first and the second dominant periods, respectively. AC_C1 and AC_C2 are the normalized autocorrelation coefficients at the first and the second dominant periods. AC_C1 and AC_DP2 were included in the feature set of the proposed method.

The autocorrelation features were found on the acceleration magnitude (Euclidean norm). For the other features, the subscript indicates on which specific input they were calculated (acceleration magnitude or one of the three acceleration signals corresponding to the reference coordinates of the accelerometer). Acceleration magnitude is insensitive to changes in the orientation of the device. In some studies, where the orientation of the device with respect to the user's body is not known in advance and can change during use, it is essential to rely only on the acceleration magnitude. In our context, however, the body positioning of a watch-like device is predictable and it is reasonable to suppose that will not vary significantly during use.

The selection of the feature set was performed starting from a larger set with about 40 features. The technique used for selection was the Correlation-based Feature Subset Selection method with greedy hill climbing search [13].

C. Personalized filter

The proposed approach to filter out non-gait instances is based on semi-supervised anomaly detection, where the training data (i.e., the personalized model) contains labeled instances only for the normal class. Each gait instance obtains an anomaly score based on Euclidean distance and nearest-neighbor analysis. The comparison between this anomaly score and a threshold enables the classification of a gait instance as normal (real gait) or abnormal (non-gait, spurious hand movements).

The set of a user's gait instances collected during the training phase is defined as $X = \{x_1, ..., x_M\}$ and *dist(a,b)* is the Euclidean distance between gait instances a and b. Let us define the distance of a gait instance g from its nearest neighbor n_g in the training set X as

$$\begin{aligned} \mathrm{dist}_{\min}(g) &= \mathrm{dist}(g, n_g), \\ n_g &= \arg\min_{i \in X} \mathrm{dist}(g, i). \end{aligned} \tag{2}$$

The mean and the standard deviation of the distances between nearest neighbors in the training set are computed:

$$\begin{aligned} \mu_X &= \frac{1}{M}\sum_{i=1}^{M} \mathrm{dist}_{\min}(x_i), \\ \sigma_X &= \sqrt{\frac{1}{M}\sum_{i=1}^{M} (\mathrm{dist}_{\min}(x_i) - \mu_X)^2}. \end{aligned} \tag{3}$$

These values are used to normalize the distance of the gait instance to be classified as explained by the following equation:

$$AS_g = \frac{\mathrm{dist}_{\min}(g) - \mu_X}{\sigma_X}. \tag{4}$$

The result of this normalization is the *anomaly score* (AS) related to the instance g. For example if the standard deviation is 0.3 and the average distance is 1.0, then an instance having the nearest neighbor at 1.6 will have $AS = 2$. Finally, AS_{th} is the threshold used to classify a gait instance as normal or abnormal, and it is selected by evaluating the trade-off between filtering out non-gait data and generating false positives (i.e., filter out real gait) as explained in the next section.

III. Experimental Evaluation

For each user, his/her own true gait instances are used to form the model, while false gait instances (produced with random hand movements) are used as non-gait examples (anomalies). The evaluation procedure consists of the following steps:

1) each normal instance of a volunteer is used to estimate the False Positive Rate (FPR) of a personalized filter trained on the remaining normal instances (leave-one-instance-out cross-validation).
2) the True Positive Rate (TPR) of the same filter is estimated with all the non-gait (abnormal) instances produced by the same volunteer.
3) FPR and TPR results are averaged over cross-validation iterations.

In the context of anomaly detection it is common to refer to anomalies as *positive* instances. In this application, a positive is thus a gait instance with non-gait data, which should be detected and discarded. Therefore, the TPR measures the proportion of non-gait instances that are identified as abnormal, whereas the FPR measures the proportion of real gait instances that are improperly identified as anomalies (in other words, a real gait instance that is filtered out represents a false positive for the personalized filter). A high TPR is key to achieve high gait detection specificity (i.e., non-gait data are discarded), whereas a low FPR is key to preserve the high detection rate (sensitivity) of the GSD technique.

For a better understanding and evaluation of the filter's performance, it is interesting to present the outcome using a ROC curve. ROC curves plot the TPR against the corresponding FPR. These two values change in relation to each other as the detection threshold varies. A typical measure used to evaluate the overall system performance using a ROC curve is the Equal Error Rate (EER), which is the point of the curve having the same (1-TPR) and FPR values.

A. Data collection

For the experiments, we used two Shimmer3 devices, embedding a TI MSP430 microcontroller (up to 24 MHz clock, 16 KB RAM, 256 KB flash). Shimmer3 is equipped with a 10 DoF inertial sensing subsystem, which includes an accelerometer, a gyroscope, a magnetometer, and a barometer. In particular, the tri-axial accelerometer is an ST Micro LSM303DLHC. One of these two Shimmer devices was worn on the wrist using a wrist band (WD), while the other was put in a front trouser pocket (PD). During the experiments, acceleration has been sampled with $\sim$ 50 Hz frequency. Samples were saved to the Shimmer's persistent memory, to ensure repeatable evaluation of methods on collected data. Traces were transferred onto a PC where they have been analyzed according to the proposed method. We also verified that the considered devices can execute the whole gait detection method in real time. In addition to acceleration, angular velocity was sampled using the embedded gyroscope – the study of these samples has been deferred to future work.

Twenty volunteers were involved in the data collection campaign (5 females, 15 males, age 26.8 ± 3.6, height 174.2 ± 8.7 cm, weight 68.7 ± 13.3 kg). Volunteers were asked to walk a corridor six times: two times at preferred pace, two times at fast pace, and two times keeping the hand closer to the sensor inside their pocket. In that way, we were able to collect (at least) three different gait patterns at two different body positions from each user. At the end of the experiment, the volunteers were also asked to perform random movements with their hands, aiming at producing fake gait detections. The experiments were video recorded to enable manual labeling of gait and non-gait periods. This was used to verify the accuracy of PD acting as supervisor.

B. Results and Discussion

We start by evaluating the performance of GSD in terms of detection rate, i.e. the proportion of real gait data that is correctly detected as part of a gait segment. The modified GSD algorithm used in this work achieves an average detection rate of 95.5%. As a baseline, we compared the results obtained on the same dataset by the algorithm described in [7]. The latter, which was designed for a device worn in a trouser pocket, achieved an average detection rate equal to 88.7%. It is important to consider the worst-case performance

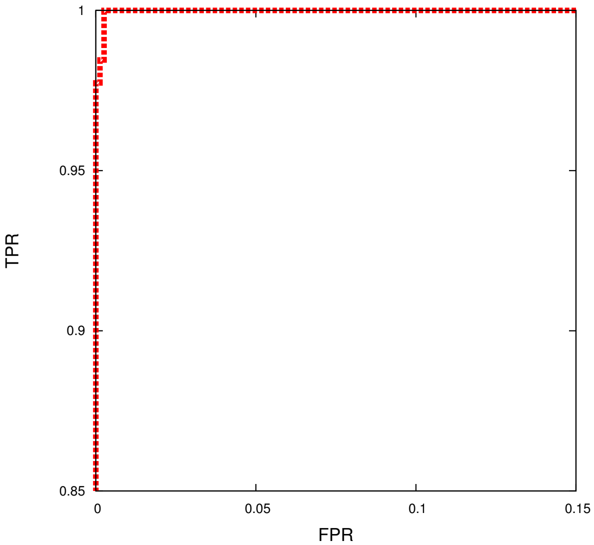

Fig. 6. ROC curve analysis

of both algorithms: the baseline algorithm showed very poor performance (below 70%) for some users, whereas the proposed technique detected about 89% of gait data in the worst case. On the other hand, this improved detection rate led to detecting at least one non-gait segment for each user in the dataset.

The core of the method is the personalized filter, which is aimed at filtering out non-gait data while preserving the high detection rate achieved by GSD. The ROC curve in Figure 6 depicts the performance of the personalized filter. The ROC shows the trade-off between TPR and FPR as the anomaly score threshold is varied. It can be observed that the filter achieves near perfect accuracy, with an EER of $\sim 0.5\%$ and an Area Under the Curve (AUC) very close to 1.

In general, the threshold value can be tuned according to the needs of the specific application for which gait detection is being used. For example, in some gait analysis applications it may be paramount to consider only real gait data – in such scenario a relatively lower threshold is preferable to maximize the filter's TPR (i.e., all non-gait instances are detected). Differently, a simple step detection application may choose the threshold maximizing the average between TPR and (1-FPR). In our experiment, this is obtained by setting the threshold to 5.4. With this setting, the filter shows perfect accuracy for 18 out of 20 users. The remaining two users have perfect TPR and a FPR lower than 3%. Thus, even for these two users the system filters out all the non-gait data, while the high detection rate of the GSD technique is nearly preserved.

IV. Conclusion and Future Work

We proposed a novel method to reliably detect gait activity with a wrist-worn accelerometer. This body position is particularly challenging, as hand movements make the identification of gait cycles more difficult. The method includes mechanisms aimed at reaching adequate sensitivity levels without compromises in terms of specificity. The evaluation was carried out with 20 subjects, who performed supervised gait experiments in a corridor. Experiments included random hand movements aimed at producing wrong detections. The results showed that the method is sound, with an average Equal Error Rate as low as 0.5%. Furthermore, the method is sufficiently lightweight to be implemented in miniaturized devices with limited resources (e.g. microcontroller-class devices with less than 16 KB of RAM). The personalized model built with the help of another device, worn during the training phase at a more favorable position, is key to obtain such results.

Future work will concern an evaluation of the method in uncontrolled environment, where users will be monitored during their habitual activities. We also plan to provide more details about real-time execution on miniaturized devices and the related energy consumption. Other aspects that need further investigation include the use of a gyroscope in addition to the accelerometer, and the development of a mechanism to automatically determine when the training phase can be concluded.

References

[1] M. E. Rosenberger, W. L. Haskell, F. Albinali, S. Mota, J. Nawyn, and S. Intille, "Estimating activity and sedentary behavior from an accelerometer on the hip or wrist," *Medicine & Science in Sports & Exercise*, vol. 45, no. 5, pp. 964–975, 2013.

[2] R. P. Troiano, J. J. McClain, R. J. Brychta, and K. Y. Chen, "Evolution of accelerometer methods for physical activity research," *British Journal of Sports Medicine*, vol. 48, no. 13, pp. 1019–1023, 2014.

[3] E. Garcia-Ceja, R. F. Brena, J. C. Carrasco-Jimenez, and L. Garrido, "Long-term activity recognition from wristwatch accelerometer data," *Sensors*, vol. 14, no. 12, pp. 22 500–22 524, 2014.

[4] M. Gjoreski, H. Gjoreski, M. Luštrek, and M. Gams, "Recognizing atomic activities with wrist-worn accelerometer using machine learning," in *Proceedings of IS International Multiconference Information Society*, Ljubljana, Slovenia, 2015, pp. 10–11.

[5] K. Ellis, J. Kerr, S. Godbole, G. Lanckriet, D. Wing, and S. Marshall, "A random forest classifier for the prediction of energy expenditure and type of physical activity from wrist and hip accelerometers," *Physiological Measurement*, vol. 35, no. 11, pp. 2191–2203, 2014.

[6] L. Atallah, B. Lo, R. King, and G. Z. Yang, "Sensor positioning for activity recognition using wearable accelerometers," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 5, no. 4, pp. 320–329, 2011.

[7] G. Cola, A. Vecchio, and M. Avvenuti, "Improving the performance of fall detection systems through walk recognition," *Journal of Ambient Intelligence and Humanized Computing*, vol. 5, no. 6, pp. 843–855, 2014.

[8] A. Brajdic and R. Harle, "Walk detection and step counting on unconstrained smartphones," in *Proceedings of ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*, Zurich, Switzerland, 2013, pp. 225–234.

[9] H. M. Thang, V. Q. Viet, N. D. Thuc, and D. Choi, "Gait identification using accelerometer on mobile phone," in *Proceedings of IEEE International Conference on Control, Automation and Information Sciences (ICCAIS)*, Saigon, Vietnam, 2012, pp. 344–348.

[10] B. Lo, S. Thiemjarus, A. Panousopoulou, and G. Z. Yang, "Bioinspired design for body sensor networks," *IEEE Signal Processing Magazine*, vol. 30, no. 1, pp. 165–170, 2013.

[11] G. Cola, M. Avvenuti, A. Vecchio, G. Z. Yang, and B. Lo, "An unsupervised approach for gait-based authentication," in *Proceedings of IEEE International Conference on Wearable and Implantable Body Sensor Networks (BSN)*, Cambridge, MA, USA, June 2015, pp. 1–6.

[12] G. Cola, M. Avvenuti, and A. Vecchio, "Real-time identification using gait pattern analysis on a standalone wearable accelerometer," *The Computer Journal*, 2017.

[13] M. A. Hall and G. Holmes, "Benchmarking attribute selection techniques for discrete class data mining," *IEEE Transactions on Knowledge and Data Engineering*, vol. 15, no. 6, pp. 1437–1447, 2003.