2014 11th International Conference on Wearable and Implantable Body Sensor Networks

Anti-Cheating: Detecting Self-Inflicted and Impersonator Cheaters for Remote Health Monitoring Systems with Wearable Sensors

Nabil Alshurafa[1], Jo-Ann Eastwood[2], Mohammad Pourhomayoun[1], Suneil Nyamathi[1], Lily Bao[1], Bobak Mortazavi[1], Majid Sarrafzadeh[1]
1. Computer Science Department, UCLA, Los Angeles, CA 90095
2. School of Nursing, UCLA, Los Angeles, CA 90095
Email: {nabil, mpourhoma, snyamathi, lilybao, bobakm, majid}@cs.ucla.edu[1], and jeastwoo@sonnet.ucla.edu[2]

***Abstract*—In remote health monitoring of patient's physical activity, ensuring correctness and authenticity of the received data is essential. Although many activity monitoring systems, devices and techniques have been developed, preventing patient cheating of an activity monitor has been a primarily unaddressed challenge across the board. Patients can manually shake an activity monitor device (sensor) with their hand and watch their physical activity points or rewards increase; we define this as "self-inflicted" cheating. A second type of cheating, "impersonator" cheating, is when subjects hand the activity sensor over to a friend or second party to wear and perform physical activity on their behalf. In this paper, we propose two novel methods based on classification algorithms to address the cheating problems. The first classification framework improves the correctness of our data by detecting self-inflicted cheatings. The second technique is an advanced classification scheme that extracts and learns unique patient-specific activity patterns from prior data collected on a patient to distinguish the true subject from an impersonator. We tested our proposed techniques on Wanda, a remote health monitoring system used in our Women's Heart Health study of 90 African American women at risk of cardiovascular disease. We were able to distinguish cheating from other physical activities such as walking and running, as well as other common activities of daily living such as driving and playing video games. The self-inflicted cheating classifier achieved an accuracy of above 90% and an AUC of 99%. The impersonator cheater framework results in an average accuracy of above 90% and an average AUC of 94%. Our results provide insight into the randomness of cheating activities, successfully detects cheaters, and attempts to build more context-aware remote activity monitors that more accurately capture patient activity.**

***Index Terms*—Cheating Detection; Wearable Body Sensors; Activity Recognition; Remote Health Monitoring System; Feature Selection;**

I. Introduction

Wearable inertial sensors are ubiquitous and becoming increasingly accurate in measuring individuals' everyday physical activity [1], [2]. Remote health monitoring (RHM) systems benefit from such sensors by learning different patterns of activity and gaining insight into a patients' daily lifestyle and behaviors, resulting in more context-aware, patient-specific interventions [3]. Recent research on people's views of physical activity research suggest that participants could be encouraged to cheat or ask a friend to wear the activity sensor on their behalf [4]. When rewards are offered based on physical activity performance, there is a risk that people, especially children, will select the easier path by cheating, where they imitate intense physical activity by manually shaking instead of wearing the activity device [5].

In our Women's Heart Health study, participants wear a smartphone on their waist that remotely monitors their physical activity. A few months into the study, nurses were wondering why some participants' health was not improving despite their activity monitors reporting intense physical activity. Our findings suggest that a few women cheated the system. Participants validated our prediction by admitting to "self-inflicted cheating," where they manually shake the phone in different directions to give the impression that they performed intense physical activity. Some participants even handed their phones to other members of the family to wear, which we call "impersonator cheating." Figure 1 depicts the two different types of cheaters we attempt to capture.

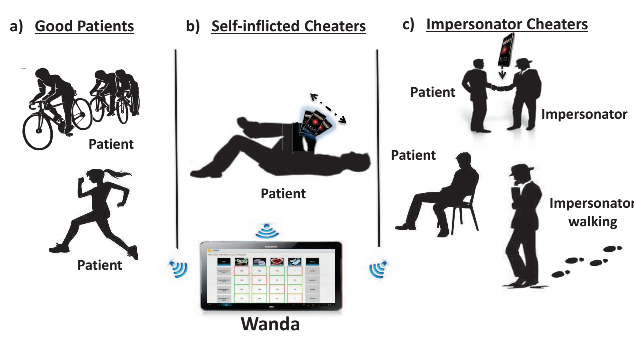

Fig. 1. a) Participants wearing their smartphone and performing physical activity. b) Self-inflicted cheating, where participants imitate intense physical activity by manually shaking the smartphone. c) Impersonator cheating, where participants hand someone else their smartphone to wear and perform physical activity on their behalf.

Thus, it is very important for an activity monitoring device to be able to distinguish cheating from real physical activity, as well as other activities of daily living. In this paper, we focus primarily on distinguishing self-inflicted cheating and impersonator cheating from other activities such as driving

978-1-4799-4959-5/14 $31.00 © 2014 IEEE
DOI 10.1109/BSN.2014.38

while wearing the smartphone, playing video games on the smartphone, running and walking. We propose a novel anti-cheating classification framework that identifies acts of self-inflicted cheating, and a novel advanced anti-cheating recognition framework to capture acts of impersonator cheating.

II. Related Work

Several activity monitoring systems have been developed to classify activity types and count steps, however, very few have attempted to detect and prevent cheating. Tomlein et. al. claims to have developed the ability to prevent cheating of an advanced pedometer application among young children by training a feedforward artificial neural network, however, accuracy of results are not provided, and the system was only compared against an unspecified intensity of walking and running [5]. Cercos and Mueller show the potential of certain participants to cheat the system; one player placed a Fitbit in the wheel of a bike to get additional amount of steps [6].

In a recent soccer exergaming platform, although the system attempts to prevent users from cheating the necessary soccer activities by detecting a certain intensity level, participants were still able to cheat the system [7]. In a prior effort, we have shown the ability to build Stochastic Approximation Models (SAM) for each specific activity type [8]. However, in this paper, our system learns patient-specific activity features to build models of activity specific to each subject, to also prevent impersonator cheating.

Several RHM systems use mobile phones to monitor physical activities [2], [3]. We have augmented our RHM system, Wanda, to distinguish cheating and other activities such as driving, playing video games, walking, and running, while at the same time learn patient-specific models to build unique patient-specific activity patterns (PSAP) and identify impersonators.

This paper is organized as follows. In Section III we describe the Wanda RHM system. Then we describe our self-inflicted and impersonation cheating classification framework in Section IV. We present our experimental setup in Section V, results in Section VI, and conclusion in Section VII.

III. Remote Health Monitoring System

In this research we used Wanda, a RHM system designed to reduce risk factors for patients. Since cheating happens frequently among patients in health monitoring systems, nurses often wonder why after three months a person supposedly performing intense physical activity does not lose weight. Wanda was designed to monitor participant's physical activity and blood pressure and collect daily/weekly questionnaires. In this research, we focus on detecting cheating vs. true physical activity, and also account for the more sedentary activities of driving and playing video games as added prevalent activities of daily living. There are several parts of the Wanda system illustrated in Figure 2. The first component is the smartphone gateway used to measure, collect and communicate data from sensors. It stores data locally as well as transmits this data through a network to a data storage center, where the raw data from multiple smartphones is collected, and data models, analytics and automated wireless messaging is performed. The proposed data analytics determine what types of activities are performed, while at the same time building activity models per subject, transmitting automated messages and alerting the nurses whenever a participant is not compliant or requires attention.

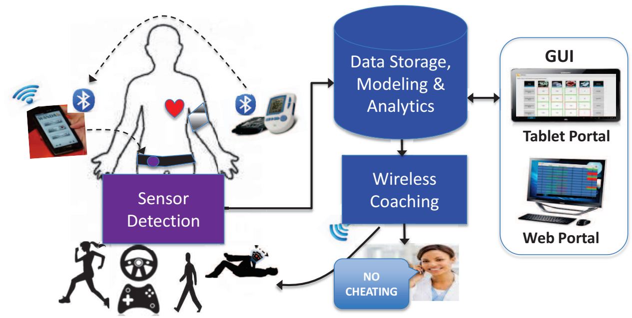

Fig. 2. Wanda system architecture designed for a Women's Heart Health Study

In order to determine the participants' physical activity level, participants are asked to wear a smartphone in a pouch around their waist. We used a metric proposed by Panasonic [9], which has been shown to have high correlation ($R^2 = 0.86$) with Doubly Labeled Water, and is one of the most accurate methods for evaluating total energy expenditure under free living conditions. K_m values, shown in Equation 1, are calculated for a given time window and are mapped to activity levels. Based on regressions performed on subjects using a metabolic cart, which estimates oxygen uptake to calculate metabolic equivalent task (MET) levels, we achieved optimal results using a 10Hz sampling rate and a time window of 5 seconds; as such, the number of samples n in 5 seconds is 50.

$$\begin{aligned} K_m &= \sqrt{\frac{1}{n-1}[Q - \frac{1}{n}(P)]}\ , \quad where \\ Q &= \sum_{i=0}^{n} x_i^2 + \sum_{i=0}^{n} y_i^2 + \sum_{i=0}^{n} z_i^2\ , \ and \\ P &= (\sum_{i=0}^{n} x_i)^2 + (\sum_{i=0}^{n} y_i)^2 + (\sum_{i=0}^{n} z_i)^2 \end{aligned} \tag{1}$$

We calculate daily physical activity levels based on the K_m values, while at the same time storing raw sensor values for further data modeling and classification. As depicted in Figure 3 from the Wanda portal, a participant admitted to cheating on specific days, but it is difficult to know simply from intensity of activity whether the individual cheated or not, hence the need for a method of cheating detection.

IV. Methodology

A. Self-inflicted Cheating Recognition

Figure 4 illustrates accelerometer signals for four types of activities performed by the same subject. The subject was asked to perform the following activities: cheat randomly with

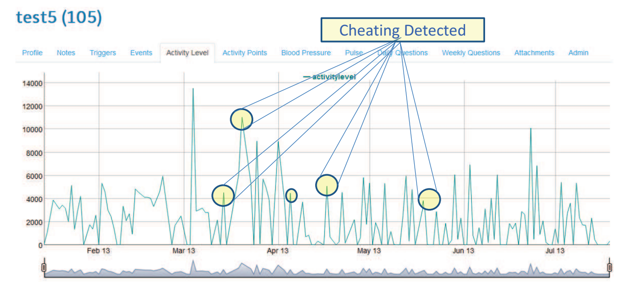

Fig. 3. Wanda portal for a subject showing several days where the subject admitted to self-inflicted cheating while laying down in bed. From daily intensity values alone it is difficult to detect cheating.

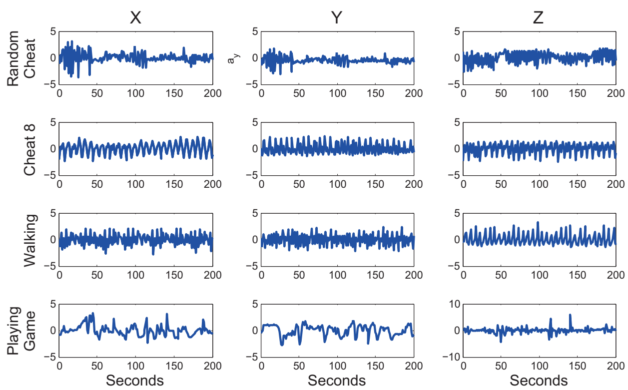

Fig. 4. Signals comparing four types of activities, two self-inflicted manual cheating and two daily activities with the smartphone in a waist-pouch. Top graph: participants randomly cheat while sitting on a chair. Second graph: participants draw a repeated "8" pattern in the air with their phone. Third and bottom graphs: participants walk and play the video game Labyrinth, which requires rotation of the smartphone to move a ball in the game. No evident pattern emerges from the random cheating.

the smartphone in their hand, cheat by repeatedly drawing an "8" pattern in the air with the smartphone, walk with the phone in the pouch around the waist, and drive with the phone in the pouch. Comparing random cheating data with that of walking and driving, it is not immediately evident if a pattern exists. However, in the proposed method, we are able to recognize cheating by extracting meaningful features that provide statistical information about the signal.

1) Feature Extraction: Several studies analyze varying features best utilized for human activity based on accelerometer data. Table I lists the main features shown to be useful in classifying activity [10], [11], [12]. Using the data from each X, Y, Z acceleration axis generates a total of 45 statistical features per segment where a segment is a fixed time subdivision of the accelerometer data. We tested segment sizes of 1, 2, 4 and 5 seconds. We achieved optimal results with a 5 second fixed time subdivision of the accelerometer data.

In order to train a classifier to detect a specific activity, feature selection is applied to properly select the features that best distinguish the activity types.

2) Feature Selection: The conventional feature selection algorithms usually focus on specific metrics to quantify the

TABLE I
FEATURE TABLE

Mean	Standard Deviation	Mean Derivatives
Median	Pairwise Correlation	Interquartile Range
Skewness	Root Mean Square	Zero Crossing Rate
Variance	Mean Crossing Rate	Kurtosis
Peak-to-Peak	Max/Min Value	Amplitude

relevance and/or redundancy to find the smallest subset of features that provides the maximum amount of useful information for prediction. Thus, the main goal of feature selection algorithms is to eliminate redundant or irrelevant features in a given feature set. Applying an effective feature selection algorithm not only decreases the computational complexity of the system by reducing dimensionality and eliminating redundancy, but also increases the performance of the classifier by deleting irrelevant and confusing information.

The two well-known feature selection categories are the filter and wrapper methods. Filter methods use a specific metric to score each individual feature (or a subset of features together), and are usually fast and much less computationally intensive. Wrapper methods usually utilize a classifier to evaluate feature subsets in an iterative manner according to their predictive power [13]. We applied the wrapper method, testing multiple combinations of feature subsets and classifiers including: kNN, BayesNet, SVM, Random Forest, and C4.5DT. The optimal feature subset and classifier combination is selected to run in real time. Figure 5 provides an illustration of the system architecture, where an optimal feature subset and classifier is trained to distinguish between cheating and other types of physical activity.

B. Advanced Impersonation Cheating Recognition

To detect an impersonator, we have to extract the main subject's activity pattern data. For this reason we acquired apriori activity data on the main subject. We built a model for each activity type using prior data collected by applying the wrapper feature selection method to find the most unique features that distinguished one person's specific activity from everyone else's. It can be challenging if a subject has similar walking patterns to his or her impersonator, but this is highly unlikely given the large variation of walking patterns between individuals [14].

As shown in Figure 6, we detected a patient-specific activity pattern (PSAP) represented by a subset of the extracted features and stored it in Wanda. Once a patient passed the self-inflicted cheating test, and a specific activity type was detected, a PSAP was requested for the detected activity type and was tested against a patient recognition classifier. If the subject passed cheating recognition tests, the system acknowledged the pureness and validity of the subjects activity data. Otherwise a nurse was informed of a potential impersonator interference.

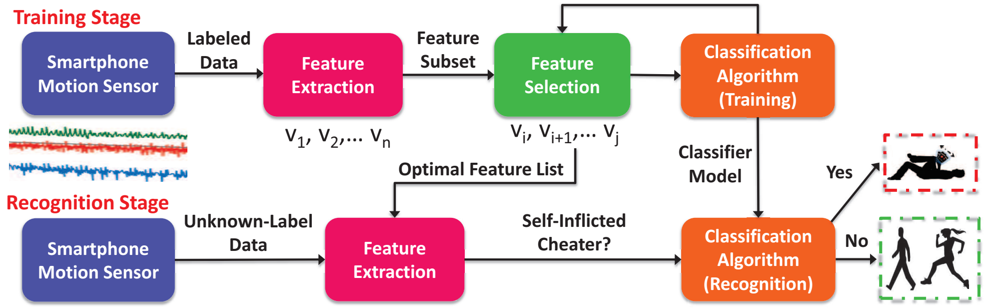

Fig. 5. Self-Inflicted cheating recognition framework.

V. Experimental Setup and Clinical Trial

Our system has been deployed in the Women's Heart Health Study, which is a UCLA IRB-approved study of 90 young black women aged 25-45 years who have a minimum of two risk factors for cardiovascular disease. In this study, 45 participants comprise the intervention group that receive nutrition and lifestyle education and smartphone-based RHM system. The control group receives usual care. The RHM transmits data in real-time using Wi-Fi and 3G/4G technology. It attempts to analyze the affects and lifestyle changes resulting from social support via automated Wireless Coaching. The intervention group receives four educational classes focused on self-management skills in nutrition, physical activity and stress-reduction using individualized goal-setting. After completing baseline screening of cholesterol levels, blood pressure, body mass index, and demographic and psychosocial questionnaires, as well as the educational classes, the women receive education on how to wear and manage the smartphones and blood pressure machines. They are also informed that the primary purpose of the smartphone is to provide a visual interface, a mechanism for automated feedback, a method to transmit blood pressure measurements, and a way to track physical activity. For this reason the women are instructed to wear the smartphone continuously in a pouch around their waist.

To test the self-inflicted cheating classification framework, we collected data from 6 subjects performing multiple activities. To test cheating, we defined 13 activity types. Five activity types were categorized as cheating activities and 8 were common activity types of daily living such as walking, running, playing video games, and driving. We requested each participant perform each activity for an average of 5 minutes. For the cheating and video game play activities the subject held the smartphone in their hand; otherwise they placed the phone in a waist pouch. The activities are defined in Table II.

The cheating activities include: random cheating, cheat-8, cheat-X, cheat-Y, and cheat-Z. For random cheating, participants were asked to sit in a chair, with the smartphone in their hand, and freely cheat the system into thinking they are physically active. They were encouraged to make us think they were performing activity, and so many attempted to shake the phone in walk-like patterns, and others shook the phone in multiple directions. In order to test a fixed pattern, we requested that each participant draw an "8" pattern in the air with the smartphone in their hand. To perform cheat-X, cheat-Y and cheat-Z, participants were requested to shake the phone randomly along each of the phone's X, Y, and Z axes. To perform the driving activity, each participant drove in their own vehicles in the city around the UCLA campus while wearing the smartphone. To perform the video game play activity, each subject sat on a chair and played Angry Birds (which requires little to no motion of the smartphone) and Labyrinth (which requires smarphone rotation). The walking and running activities were tested with varying intensity levels as defined in Table II on a Merit Fitness 715T Plus treadmill. To test the self-inflicted cheating on each subject, we performed Leave One Out Cross Validation. We created 5 categories of activities, and ensured that we randomly selected an equal number of training samples for each category to eliminate bias.

TABLE II
Activity Types

C1	Cheat Random		W1	Walking 2.5mph
C2	Cheat-8		W2	Walking 3.5mph
C3	Cheat-X		W3	Random Walk
C4	Cheat-Y		R1	Running 4mph
C5	Cheat-Z		R2	Running 5mph
P1	Play Game (Angry Birds)		D1	Driving
P2	Play Game (Labyrinth)			

To test the impersonator cheating classification framework, we focused primarily on the walking and running activities, because physical activity is what matters in the given study. Again, we selected 6 subjects and asked them to perform running and walking activities for five minutes, while at the same time selecting other remaining subjects to perform the same activities on their behalf. The testing data set includes samples from subject X as well as data samples for the same activity collected from other subjects. To test the impersonator recognition framework, we tested the classifiers' ability to distinguish one subject's activity from another's (e.g. Subject 1 walking vs. another's walking) and analyze the classifiers' accuracy and receiver operating characteristic (ROC) area under the curve (AUC) using each subjects PSAP. We repeated this experiment for each subject.

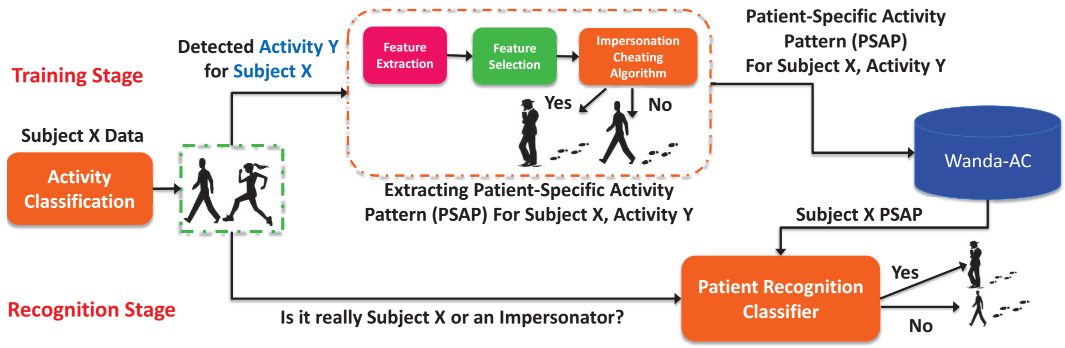

Fig. 6. Impersonator cheating recognition framework.

VI. Results and Description

A. Self-inflicted Cheating

According to our classification results, acts of self-inflicted cheating can be detected with high accuracy. The classifier that yielded the best result was the Random Forest classifier. Table III provides a confusion matrix for the outcome of the self-inflicted classification framework. Results show a 91.9% cheating recall. The precision for cheating is 83.2% because running samples are sometimes misclassified as cheating. Running activity types result in 81.8% recall, with misclassifications resulting from cheating samples. Walking, on the other hand, yields 93.2% and 91.8% recall and precision, respectively. Some walking was misclassified as playing video games, and that is probably due in part to playing the Labyrinth video game, which requires rotating the smartphone to move the ball in motions that could at times exhibit features similar to walking.

While our primary focus is cheating and physical activity, it is also interesting to note from Table III that driving resulted in 77.0% recall and 98.5% precision. This is interesting because with such high precision, this informs us that when our classifier detects a subject driving, we are very certain they are driving, but with lower recall, some driving activities could be missed. Playing video games resulted in a high recall of 99.2% and a precision of 81.2%. It is interesting to note that a lot of driving was confused with playing video games, which is justified by the fact that drivers must wait at traffic light signals exhibiting little or no motion in the hip; similarly, playing Angry Birds also requires little to no smartphone motion.

Table IV provides the accuracy and area under the curve of the self-inflicted classifier for each Leave One Out Cross Validation test run. The table shows a high average accuracy of 90% and a high discriminating power with an average area under the ROC curve of 99%. It is interesting to note that Subjects 3 and 5 seemed to exhibit more features similar to those of the other subjects (used in training), resulting in higher average accuracy than others. However, Subject 4 seemed to be a bit more unique in performing the defined activities, resulting in lower classification accuracy, suggesting that she was not as well represented by other subjects in the dataset. Our results show promise in the ability to build models of patients that are similar to each other.

TABLE III
Confusion Matrix of Runs for the Self-Inflicted Random Forest Classifier

	Predicted Outcome					
Activity	Cheat	Drive	Game	Run	Walk	Recall
Cheat	237	1	4	14	2	91.9%
Drive	0	194	40	0	18	77.0%
Game	0	2	250	0	0	99.2%
Run	47	0	0	216	1	81.8%
Walk	1	0	14	2	234	93.2%
Precision	83.2%	98.5%	81.2%	93.1%	91.8%	

TABLE IV
Accuracy and AUC of the Leave One Out Cross Validation of the Self-Inflicted Classifier.

Leave Subject X Out	Accuracy	AUC
Subject 1	86%	100%
Subject 2	90%	99%
Subject 3	97%	100%
Subject 4	82%	96%
Subject 5	96%	100%
Subject 6	91%	100%
Average	90%	99%

B. Impersonator Cheating

If a sample is classified as walking and running, what is to say it is the intended subject that is walking or running as opposed to an impersonator? Dealing with the second type of cheater is very challenging and requires learning. The Random Forest classifier also outperformed other classification schemes in detecting the acts of impersonator cheating. We tested the impersonator cheating recognition and the results in Table V show the high performance of the proposed method to distinguish a specific subjects walk from that of other subjects. As can be seen, Subject 4 is the most distinguishable subject in terms of walking and running, resulting in a high walking and running accuracy of 100% and 97%, respectively. The subject with the least distinguishability was Subject 5, with an accuracy of 81% in walking, and 85% in running. However, the average walking accuracy of both running and walking is

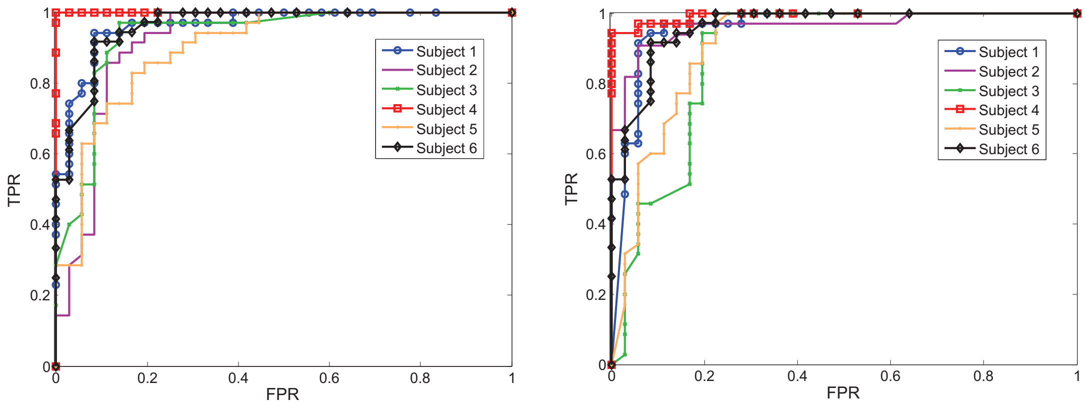

(a) Walking activity

(b) Running activity

Fig. 7. Impersonation cheating classifier ROC curve.

TABLE V
ACCURACY AND AUC OF THE IMPERSONATOR CLASSIFIER, WHICH DISCRIMINATES SUBJECT X'S ACTIVITY FROM ALL OTHERS' ACTIVITIES.

Subject X	Walk ACC	Walk AUC	Run ACC	Run AUC
Subject 1	93%	96%	93%	96%
Subject 2	87%	92%	93%	96%
Subject 3	90%	93%	87%	89%
Subject 4	100%	100%	97%	99%
Subject 5	81%	90%	85%	91%
Subject 6	89%	96%	85%	93%
Average	90%	94.5%	90%	94%

90%, and the average AUC for both running and walking is about 94%.

Figure 7 also shows results for each subject's learned PSAP in combination with the Random Forest classifier. The results for each subject are all above 90% AUC, which shows excellent discrimination in correctly classifying impersonators.

VII. CONCLUSION

In order to arrive at the proper inferences from our physical activity monitoring systems, we must be able to validate the correctness and authenticity of such devices. In this paper, we proposed a novel method based on classification and machine learning techniques that enhances correctness of our remote monitoring of physical activity by distinguishing acts of self-inflicted cheating from physical activity and other activities of daily living, yielding an accuracy of 90% and an excellent ROC curve AUC of 99%. We have also developed a novel classification-based method that learns patient-specific activity patterns to enhance the authenticity of the data by distinguishing acts of impersonator cheating in activities of physical activity, resulting in an average accuracy of 90% and a ROC curve AUC of 94%. Our work detected acts of self-inflicted and impersonator cheating in the Women's Heart Health study, and contributed to purifying our study results and data while deterring participants from cheating, shifting the subjects focus from cheating to performing true physical activity.

REFERENCES

[1] P. S. Freedson, K. Lyden, S. Kozey-Keadle, and J. Staudenmayer, "Evaluation of artificial neural network algorithms for predicting METs and activity type from accelerometer data: validation on an independent sample," *J. Appl. Physiol.*, vol. 111, no. 6, pp. 1804–1812, Dec 2011.

[2] S. Patel, H. Park, P. Bonato, L. Chan, and M. Rodgers, "A review of wearable sensors and systems with application in rehabilitation," *Journal of NeuroEng. and Rehabilitation*, vol. 9, no. 1, pp. 1–17, 2012.

[3] H. Ketabdar and M. Lyra, "System and methodology for using mobile phones in live remote monitoring of physical activities," in *Technology and Society (ISTAS), 2010 IEEE Int. Symp. on*, June 2010, pp. 350–356.

[4] J. Kirby, C. Tibbins, C. Callens, B. Lang, M. Thorogood, W. Tigbe, and W. Robertson, "Young People's Views on Accelerometer Use in Physical Activity Research: Findings from a User Involvement Investigation," *ISRN Obes*, vol. 2012, p. 948504, 2012.

[5] M. Tomlein, P. Bielik, P. Krátky, M. Barla, and M. Bieliková, "Advanced pedometer for smartphone-based activity tracking." in *Proceedings of the 6th International Conference on Health Informatics*, 2012, pp. 401–404.

[6] R. Cercos and F. F. Mueller, "Watch your steps: Designing a semi-public display to promote physical activity," in *Proceedings of The 9th Australasian Conference on Interactive Entertainment: Matters of Life and Death*, ser. IE '13. New York, NY, USA: ACM, 2013, pp. 2:1–2:6.

[7] B. Mortazavi, K. C. Chu, X. Li, J. Tai, S. Kotekar, and M. Sarrafzadeh, "Near-realistic motion video games with enforced activity," in *BSN*, 2012, pp. 28–33.

[8] N. Alshurafa, W. Xu, J. J. Liu, M.-C. Huang, B. Mortazavi, M. Sarrafzadeh, and C. Roberts, "Robust human intensity-varying activity recognition using stochastic approximation in wearable sensors." in *BSN*. IEEE, 2013, pp. 1–6.

[9] Y. Yamada, K. Yokoyama, R. Noriyasu, T. Osaki, T. Adac hi, A. Itoi, Y. Naito, T. Morimoto, M. Kimura, and S. Oda, "Light-intensity activities are important for estimating physical activity energy expenditure using uniaxial and triaxial accelerometers," *Eur. J. Appl. Physiol.*, vol. 105, no. 1, pp. 141–152, Jan 2009.

[10] N. Ravi, N. D, P. Mysore, and M. L. Littman, "Activity recognition from accelerometer data," in *Proceedings of the 17th Conference on Innovative Applications of AI (IAAI.* AAAI Press, 2005, pp. 1541–1546.

[11] M. Zhang and A. A. Sawchuk, "A feature selection-based framework for human activity recognition using wearable multimodal sensors," in *Proceedings of the 6th Int. Conf. on Body Area Networks*, ser. BodyNets '11, 2011, pp. 92–98.

[12] W. Xu, M. Zhang, A. Sawchuk, and M. Sarrafzadeh, "Robust human activity and sensor location corecognition via sparse signal representation," *Biomedical Engineering, IEEE Transactions on*, vol. 59, no. 11, pp. 3169–3176, 2012.

[13] I. Guyon and A. Elisseeff, "An introduction to variable and feature selection," *J. Mach. Learn. Res.*, vol. 3, pp. 1157–1182, Mar. 2003.

[14] D. A. Winter, "Kinematic and kinetic patterns in human gait: Variability and compensating effects," *Human Movement Science*, vol. 3, no. 12, pp. 51 – 76, 1984.