Pervasive and Mobile Computing 38 (2017) 281–295

Contents lists available at ScienceDirect

Pervasive and Mobile Computing

journal homepage: www.elsevier.com/locate/pmc

Position-aware activity recognition with wearable devices[☆]

Timo Sztyler[a,*], Heiner Stuckenschmidt[a], Wolfgang Petrich[b]

[a] University of Mannheim, Germany

[b] Roche Diabetes Care GmbH, Mannheim, Germany

ARTICLE INFO

Article history:
Available online 30 January 2017

Keywords:
Pervasive computing
Activity recognition
Wearable sensors
Machine learning
On-body position detection

ABSTRACT

Reliable human activity recognition with wearable devices enables the development of human-centric pervasive applications. We aim to develop a robust wearable-based activity recognition system for real life situations where the device position is up to the user or where a user is unable to collect initial training data. Consequently, in this work we focus on the problem of recognizing the on-body position of the wearable device ensued by comprehensive experiments concerning subject-specific and cross-subjects activity recognition approaches that rely on acceleration data. We introduce a device localization method that predicts the on-body position with an F-measure of 89% and a cross-subjects activity recognition approach that considers common physical characteristics. In this context, we present a real world data set that has been collected from 15 participants for 8 common activities where they carried 7 wearable devices in different on-body positions. Our results show that the detection of the device position consistently improves the result of activity recognition for common activities. Regarding cross-subjects models, we identified the waist as the most suitable device location at which the acceleration patterns for the same activity across several people are most similar. In this context, our results provide evidence for the reliability of physical characteristics based cross-subjects models.

1. Introduction

Activity Recognition is an active field of research in pervasive computing [1–3]. The development of wearable devices such as smart-phones, smart-watches, and fitness wristbands feature a variety of sensors that are carried all day long by many people (compare [2]) and provide new opportunities for continuous monitoring of human activities such as running, sitting, standing, or walking [4]. A problem of many existing studies on the subject is that they are conducted in a highly controlled environment. In consequence, the results of these studies often do not carry over to real world applications. Our aim is to develop robust activity recognition methods based on mobile device sensors that generate high quality results in a real world setting.

At present, activity recognition in a real world scenario goes along with several unaddressed problems. First, commonly it is up to the user where the wearable device is worn and often the device position is chosen with regard of the performed activity. This means that also transitions between positions have to be detected. Further, most of the existing approaches focus on a subject-specific approach, i.e., the target user has to collect and label data. This is often not feasible, especially

☆ A preliminary version of this work appeared in the Proceedings of the IEEE International Conference on Pervasive Computing (PerCom 2016).

* Corresponding author.
E-mail addresses: timo@informatik.uni-mannheim.de (T. Sztyler), heiner@informatik.uni-mannheim.de (H. Stuckenschmidt), wolfgang.petrich@roche.com (W. Petrich).

http://dx.doi.org/10.1016/j.pmcj.2017.01.008
1574-1192/

in a healthcare scenario where elders or patients should be observed and are unable to perform all activities for a certain amount of time to collect enough training data. The complement could be a cross-subjects approach that relies on labeled data of certain people except the target user. More precisely, a classification model is trained on data of several people to classify the activities of another unseen person. Hence, training and testing data are disjoint sets and these sets represent different people. In this context, the idea is to take advantage of similar body movement patterns across people.

Preceding works have already state important insights concerning these problems. These include which body areas provide different information for the same activity which in turn is essential for using sensor data across people. In this context, the relevant on-body positions that should be distinguished are head, upper arm, forearm, chest/waist, thigh, and shin [5]. Further, they stated that the on-body device position has an influence on the quality of activity recognition. Focusing on different cross-subjects approaches, most of the existing works examined leave-one-subject-out [6,7], i.e., the consideration of all available labeled data except the data of the target person. This often performs significantly worse than a subject-specific approach due to different acceleration patterns, e.g., an elder has different locomotion patterns than a child. However, cross-subjects models can be personalized by co-training [8] or adaptation of parameters [9] but a satisfiable initial recognition model is preferable. Thus, the better the initial model performs the lower the costs for the personalization [10]. In general, these works rely on acceleration sensors and also provide evidence for its reliability [4]. Moreover, an acceleration sensor has a low power consumption which makes it interesting concerning continuous sensing over a complete day.

This paper is an extension of a previous work [11]. We already focused on the detection of relevant on-body positions in the context of everyday movements and activities using a single accelerator sensor. Further, we examined the impact of the recognized position information on the accuracy of the activity recognition. As an extension but not limited to, we also focus on cross-subjects approaches especially concerning the different device on-body positions. We aim to investigate the acceleration patterns of our subjects to identify groups which enable to build more reliable initial cross-subjects models for people that are unable to collect and label data for a subject-specific approach.

The main contributions of our work are the following:

- We show that random forest based approaches are able to recognize the device on-body position (89%) and performed activity concerning subject-specific (84%) and cross-subjects (79%) approaches.
- We show that transferring labeled data between people of the same gender and with a similar level of fitness and statue is feasible for cross-subjects activity recognition for people that are unable to collect required data.
- We present comprehensive experiments concerning subject-specific and cross-subjects activity recognition in consideration of all relevant on-body device position and physical activities.

The paper is structured as follows: In Section 2, the related work of the focused research questions is summarized. Then, we outline the data collection phase and present our data set. Section 4 covers the process of feature extraction and classification where we introduce our approach for on-body position detection, position-aware activity recognition, and cross-subjects activity recognition. In Section 5, we present our experimental results. Finally, Section 6 covers the conclusion and future work of this paper.

2. Related work

Acceleration sensor based activity recognition has been studied for many years [4,12,13] and enables to recognize common physical activities such as walking or running. The spreading of wearable devices furthered the feasibility of this approach in the real world but also gives priority to less focused issues. These include cross-subjects based activity recognition models and the varying on-body position of the wearable devices.

The cross-subjects activity recognition problem got significantly less attention than subject-specific activity recognition approaches [14,6] while especially often elders and patients are unable to collect required training data. Commonly, researchers focus on leave-one-subject-out [15] where most researchers state that the models perform worse. Their results show that especially acceleration patterns of dynamic activities (e.g., climbing stairs) may differ across different users [10]. This result can be attributed to the different physical characteristics, e.g., a child walks faster than an elder or a woman could have a different body movement than a man. Several researchers already hypothesized that physical characteristics, i.e., gender, weight, height, and fitness level, could be reliable indicators to choose specific people for a cross-subjects model [16,14]. However, focusing on common physical activities, this is still an open issue. Besides, researchers also evaluated a pairwise approach [17], i.e., trained the model on data of one person and evaluated the performance on another. However, they state that this approach often cannot yield accurate results. Moreover, the performances of these kinds of approaches concerning different on-body device positions are also still unclear.

To avoid the effort that goes along with collecting and labeling training data, researchers have also focused on creating a cross-subjects model and applying personalization. The idea is to adapt an initial cross-subjects model to the behavior of the target user to improve the recognition rate by relying on unlabeled but classified acceleration data. For the purpose of personalization, they focused on co-training [8], parameter adaption [17], and incremental learning [18]. Their result show that personalization is feasible and that a cross-subjects model can improve in a short time. However, their results also indicate that if the initial model performs worse the personalization may go along with more effort [10]. For that reason, a cross-subjects approach except leave-one-subject-out is preferable especially if the available labeled data set covers certain

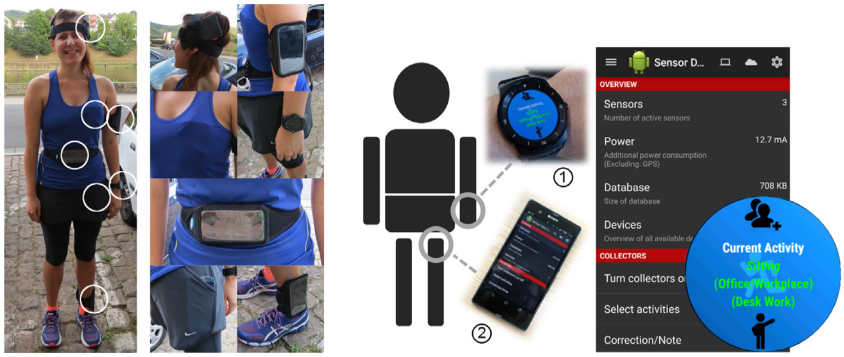

(a) Sensor placement. (b) Sensor data collector.

Fig. 1. The framework consists of a wear (1) and hand (2) app (right) which enable to record each sensor and also provide labeling functions. The subject wears (left) the wearable devices on the head, chest, upper arm, wait, forearm, thigh and shin (top down).

amount of different people. In this context, Vo et al. [10] also state that an increasing number of people goes along with a decreasing activity recognition rate.

Focusing on the on-body device position recognition shows that this problem got also less attention. Initially, researchers investigated position-independent activity recognition [19] but concluded that the recognition rate can differ significantly. Subsequently, researchers conducted experiments concerning the device position and its influence on the recognition rate and stated that this information increases the accuracy of the activity recognition but the opinions regarding the impact are divided [20,13,1]. Reviewing their results, points out that the different statements may result from the sets of positions and activities that were considered. Indeed, so far nobody considered all relevant on-body positions in context of common physical activities in one study. Therefore, it is still unclear how accurate each relevant position can be detected regarding different activities.

So far, the localization problem was only addressed by a couple of researchers. Kunze et al. published one of the first approaches that first tries to detect if the subject is walking and then used specific patterns of sensor readings to derive the current device position [21]. However, this approach is limited due to the small set of selected positions and the fact that position changes are not recognized if the subject does not walk. Deriving the device positions hand, bag, or pocket directly from the performed physical activity has shown that the effect of the location information on the accuracy of the activity recognition depends on the performed activity [13].

In general, due to the importance of the sensor placement in context of activity recognition, several researchers also investigated the influence and effect of different positions concerning the performed activities [20,5,1]. The results of these studies show that there are seven different body locations that behave differently in activity recognition, i.e., forearm, head, shin, thigh, upper arm, and waist/chest. Dividing these body parts (e.g., head or shin) into smaller regions does not improve the accuracy [5]. In addition, further studies have shown that the optimal sensor placement depends on the activity that has to be recognized [20]. As a result, the benefiting of the position information and also the feasibility to derive device positions by an accelerometer is stated.

3. Data set

In this paper, we investigate the detection of the on-body position of a wearable device, its influence on the quality of activity recognition, and the feasibility of cross-subjects activity recognition. For this purpose, we created a data set[1] which covers, among others, the acceleration data of the activities climbing stairs down (A_1) and up (A_2), jumping (A_3), lying (A_4), standing (A_5), sitting (A_6), running/jogging (A_7), and walking (A_8) of fifteen subjects (age 31.9 ± 12.4, height 173.1 ± 6.9, weight 74.1 ± 13.8, eight males and seven females). For each activity, we recorded simultaneously the acceleration of the chest (P_1), forearm (P_2), head (P_3), shin (P_4), thigh (P_5), upper arm (P_6), and waist (P_7). Each subject performed each activity roughly 10 min except for jumping due to the physical exertion ($\sim$1.7 min). In detail, we recorded for each position and axes 1065 min of sensor data which is equally distributed between male and female subjects. Concerning cross-subjects activity recognition, we want to emphasize that our group of subjects comprises several different kind of people. Hence, there are significant differences concerning fitness, physique, age, weight, height, and movement behavior. To the best of our knowledge the result is the most complete, realistic, and transparent data set for on-body position dependent approaches that is currently available.

The required data was collected using customary smart-phones and a smart-watch[2] which were attached to the mentioned positions (see Fig. 1(a)). The devices were synchronized with the time service of the network provider and the

1 http://sensor.informatik.uni-mannheim.de.
2 “Samsung Galaxy S4” and “LG G Watch R”.

Table 1
Summary of considered feature methods.

	Methods
Time	Correlation coefficient (Pearson), entropy (Shannon), gravity (roll, pitch), mean, mean absolute deviation, interquartile range (type R-5), kurtosis, median, standard deviation, variance
Frequency	Energy (Fourier, Parseval), entropy (Fourier, Shannon), DC mean (Fourier)

accelerometer was sensed with a sampling rate of 50 Hz where the data was stored on a local SD card. The sampling rate was chosen with consideration of battery life and with reference to previous studies [17,22]. The recording of the data was performed using a self-developed sensor data collector and labeling framework. The framework consists of a *Wear* (1) and *Hand* (2) application (see Fig. 1(b)) which interact with each other via Bluetooth. The application provides the possibility to control the built-in sensors, to specify the sampling rate, and to record several sensors simultaneously. The binary[3] and the source code[4] of this application are publicly available.

To attach the devices to the relevant body positions, common objects and clothes were used such as a sport armband case, trouser pocket, shirt pocket, or the bra. There was no further fixation of the devices to closely resemble their use in everyday life. We used a belt to attach a phone to the head to avoid that the subject had to hold this device during the performance of the activities. This simulates that the subject is making a phone call.

The data collection took place under realistic conditions, i.e., the subjects walked through the city, jogged in a forest, or climbed up the stairs of a guard tower of an old castle. The order of the activities was left to the subjects but they were instructed to stand idle for a few seconds before and after the activity was performed. Concerning the activities, there were no instructions. It was up to the subjects, e.g., how fast they wanted to walk or how they wanted to sit. In this context, typically the subjects used their smart-phone, talked with somebody else, or were eating and drinking something.

Each movement was recorded by a video camera to facilitate the usage of our data set also by other people. Our data set is available[1] and covers beside the mentioned acceleration data also GPS, gyroscope, light, magnetic field, and sound level data which were also recorded during the data collection phase but is not considered in the following. Besides, there is also a detailed description of each subject including images of the attached devices and a short report.

4. Method

Following most existing works, we use a supervised approach, both for on-body localization and for activity recognition. The introduced data set was used as training data and for evaluation. In the following, we describe the features generated from the sensor data and the learning methods and strategies used in our study.

4.1. Feature extraction

The essential idea behind generating features from time dependent data streams is to segment the recorded data into windows and compute a feature vector for each window. Preceding studies in the context of activity recognition already examined different settings regarding the window size and meaningful features [14,23]. They state that overlapping windows are more suitable because they can handle transitions more accurately. Further, the window size depends on the kind of activities which should be recognized. In our context, most of the existing studies considered a size between one and three seconds [19,3,24]. However, so far there is no agreed set of features. Indeed, a comparison of the different but overlapping feature sets of previous studies is difficult due to their different settings and goals. Nevertheless, some researchers have compared different groups of features and also state that frequency-based features improve the accuracy of the recognition [24].

Based on these results, we use windows which overlap by half and have a length of one second. Further, we consider the most common time- and frequency-based features that were used in previous work (see Table 1) where time-based feature values are transformed into frequency-based feature values by applying *Discrete Fourier transform*[5].

We also computed gravity-based features that provide information of the device orientation. The gravity component was extracted from the recorded accelerometer data. We applied a *low-pass filter*[6] to separate the acceleration and gravitational forces to derive the gravity vectors. These vectors enable to determine the orientation of the device by computing the angles between them also known as *roll* and *pitch* (see Fig. 2). The *azimuth* angle, however, cannot be calculated because the direction of north is required. This means that it is not possible to derive if the device is back-to-front. Further, we only consider absolute values of the acceleration so that we do not distinguish if the device is upside down. We consider these

3 https://play.google.com/store/apps/details?id=de.unima.ar.collector.

4 https://github.com/sztyler/sensordatacollector.

5 The Fourier transformation can be applied with different scaling factors. We use the *JTransforms* implementation (https://github.com/wendykierp/JTransforms) which scales by 1.

6 A low-pass filter passes values which have a lower frequency as the specified *cutoff frequency* and attenuates values that have a higher frequency.

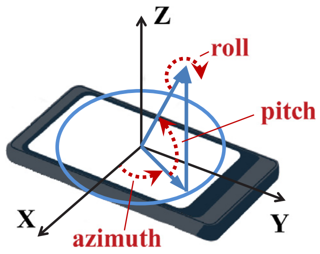

Fig. 2. Gravity-based features. The coordinate system is defined in reference to the screen. The acceleration of the device is measured along the axes. The gravity enables to compute the angle between the axes to determine the orientation (roll, pitch). To calculate azimuth, the direction of north is required.

four cases as the same orientation. To be more flexible and avoid overfitting, we also transform the roll and pitch angles in one of sixteen predefined discretized orientations. Besides, the gravity-based features are only considered in the context of on-body position detection or static activities. Previous works already showed that this feature is beneficial to distinguish between standing and lying.

The feature extraction process was performed with a self-developed framework that computes all mentioned features. The framework is available[7] and enables to specify the mentioned settings. As a result, the framework returns a list of feature vectors which are in the following further processed.

4.2. Random forest classifiers

Decision trees have already successfully been used for activity recognition, however, it is well known that classical decision trees are sensitive to overfitting when the generated trees become very deep. In order to overcome the overfitting problem, ensemble methods have been proposed that balance the results of multiple decision trees that have been trained on different parts of the training data. Random forest classifiers are one of these ensemble methods that have been proposed by Breiman [25]. A random forest classifier is constructed in the following way:

Let $D = \{(\overline{x}_1, y_1), \dots, (\overline{x}_n, y_n)\}$ be a learning problem with feature vectors $\overline{x}_i$ and results y_i. In a first step a number of samples $S_1, \dots, S_m$ are drawn from D using sampling with replacement. For each sample S_i, a decision tree classifier f_i is trained using a variation of the classical decision tree learning algorithm that uses feature bagging. This means that for each branching decision in the decision tree construction only a randomly selected subset of feature vectors is taken into account. This is necessary to ensure that the different generated decision trees are uncorrelated [26]. In this context, the decision tree considers the information gain of each feature to determine the importance during the construction.

The resulting set of uncorrelated decision trees can now be used to determine the outcome for an unseen feature vector $\overline{x}'$ based on the principle of bagging. In particular, the result is determined by averaging over the predicted results of all individual decision trees as follows:

$$\hat{f}(\overline{x}') = \frac{1}{n}\sum_{i=1}^{n} f_i(\overline{x}'). \quad (1)$$

For the case of a classification problem, the combined classifier essentially performs a majority vote over the outcomes of the individual decision trees. It has been shown that bagging prevents the overfitting problem as the combination of multiple classifiers has a significantly lower variance than an individual classifier.

4.3. On-body position detection

We treat position detection as a multi-class classification problem with target classes being head, upper arm, forearm, chest, waist, thigh, and shin that correspond to the relevant positions according to Vahdatpour et al. [5].

In initial experiments, we observed a major problem when trying to distinguish between different device positions while considering all performed activities. More precisely, data of the activities lying, standing, and sitting frequently leads to misclassification of device positions. This is caused by the fact that in context of these three activities the human body only has a slight acceleration so that the computed feature vectors are not easily distinguishable. To address this problem, we distinguish between static (standing, sitting, lying) and dynamic (climbing up/down, jumping, running, walking) activities and consider these two groups in the following as two types of activity-levels. This enables to consider different features sets. Hence, we train a classifier that distinguishes between static and dynamic activities that is used as a first step in the position detection process (see Fig. 3). A similar distinction has been made by Yang et al. [27] to improve the accuracy of activity recognition.

7 https://github.com/sztyler/sensorfeatureextraction.

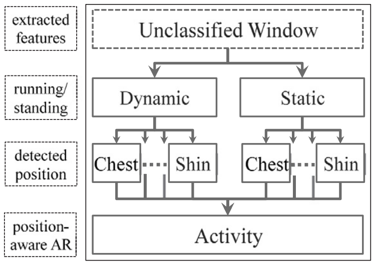

Fig. 3. Activity Recognition. The nodes illustrate the target class and the edges the applied classifier. The current window is classified as "dynamic" (climbing, jumping, running, walking) or "static" (standing, sitting, lying). Subsequently, the device position is recognized and a position specific classifier applied to derive the performed activity.

We trained both classifiers using stratified sampling combined with 10-fold cross validation to ensure that all folds cover the same ratio of classes. Further, to make the result more stable, we performed 10 runs where each time the data set was randomized and the ten folds were recreated. The classifiers were trained and evaluated for each subject individually. Thus, we did not consider several subjects at once.

4.4. Subject-specific activity recognition

In the activity recognition phase, we aim to detect the activities climbing stairs up and down, jumping, lying, running, sitting, standing, and walking. In this context, we evaluate the impact of the information of the device position. For this purpose, we construct position-independent and position-aware activity classifiers and compare their performance on our data set.

The *position-independent* activity recognition approach simply consists of a single classifier per subject that is trained on all data independent of the recognized device position. We expect this recognition approach to perform sub-optimal as the motion information from the sensors can be assumed to be very different in the different positions for the same activity.

The *position-aware* activity recognition approach consists of a set of individual classifiers, one for each device position and each subject. The classifier to be used is determined in a position recognition step that is executed before the actual activity recognition. Fig. 3 provides an overview of the detection process: First, the unlabeled record is classified as a dynamic or a static activity. As mentioned above, this step is necessary as we can more reliably detect the device position if we know whether the performed activity is a static or a dynamic activity. Then, the position of the device is recognized with an activity-level dependent classifier that uses a feature set that has been optimized for the type of activity. Finally, the performed activity is recognized by selecting and applying the classifier for the detected device position. Obviously, the performance of the position-aware activity recognition approach relies on the correct identification of the device position. Therefore, to test the feasibility of this approach, we use the results of the activity-level dependent position detection experiments – including all mistakes made – as input for the activity recognition experiments.

With the availability of wearable devices different from mobile phones – e.g. smart-watches and smart-bands – a natural question that also arises is whether using these devices in addition to a smart-phone can further improve the recognition rate. Hence, in addition we also address this question by investigating activity classification based on sensor information from multiple sensors. In order to be as general as possible, we consider combinations of two and three acceleration sensors at arbitrary positions. We consider the same set of features as before, but compute them separately for each sensor. Therefore, for each considered on-body position we have separate features. Subsequently, the resulting feature vectors of the windows that describe the same point in time are unified in a single feature vector.

4.5. Cross-subjects activity recognition

The initial idea of a cross-subjects based model is to perform activity recognition also for people, e.g. elders, which are unable to collect and label required data but, e.g., need to be observed. Commonly, a cross-subjects approach relies on labeled sensor data of several people where the most known approach is leave-one-subject-out. Thus, a single classifier is trained on all available labeled data expect data of the target person. Compared to our subject-specific approach, we focus on the performance of different cross-subjects approaches depending on the individual on-body device positions, i.e., we assume that we know the device position. However, we also evaluate how well the positions themselves could be recognized. For that purpose, we construct and evaluate the following cross-subjects approaches: *Randomly*, *Leave-One-Subject-Out*, *Top-Pairs*, and *Physical*. Especially, the physical-based approach could be promising as this idea was already hypothesized but not investigated in several previous works. For all approaches, we follow a group-based approach where the groups are dynamically determined and can overlap for different subjects. Thus, a group represents certain people whose labeled data is considered to train a classification model for an unseen subject.

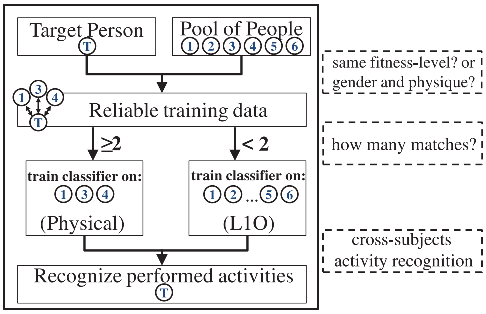

Fig. 4. Cross-subjects activity recognition by relying on demographic characteristics, i.e., fitness, gender, and physique. For instance, to determine the activities of the target person T, we do this based on the known labeled data of subjects 1, 3, and 4 (matches) which have the same fitness-level or gender and physique as T.

Leave-One-Subject-Out This approach was most often considered in related works and performs often differently depending on the considered data set. We build for each subject a classifier that relies on all available labeled data except the target person. We consider this approach as baseline.

Top-Pairs We compare our subjects pairwise to identify the best matches for each subject, i.e., we trained a classifier on data of one person and evaluated the performance on another. Based on these results, we build a classifier for a target user that consists of the top five matches. In this context, it is unclear if the best matches taken together perform better or even worse due to contradictions. Indeed, this approach can only be evaluated if labeled data of the target person is available. For that purpose and in reference to our scenario, we consider only one minute per activity of the available labeled data of the target user.

Physical In initial experiments, we investigated whether demographic characteristics, in our case gender, fitness, and physique can be used to determine a group of people whose data can be used to recognize activities of a previous unseen subject. For this purpose, we identified these characteristics for each subject from the data set. While gender and physique (strong and slim) were determined based on the videos of the exercises, we took the distance covered in 10 min running to classify the subjects into five fitness levels. The choice based on the idea that people with the same fitness level have similar patterns concerning running while the gender and physique could be characterizing for walking. For clarification, Fig. 4 illustrates this process. In case that there is at most one match, we fallback and apply *leave-one-subject-out* due to results of pairwise approaches of related work [17].

Randomly As an additional reference, we also build classifiers where the number of considered people and also the people themselves are chosen at random except the target user. We repeat this approach ten times and consider the average as recognition rate.

During our experiments, we initially focus on dynamic activities because we believe that the acceleration patterns of static activities are less characterized by the individual behavior. We examine the performance and benefits of the introduced cross-subjects models but also the individual performance in context of each on-body position. Finally, we discuss and compare the results of our subject-specific and cross-subjects approaches also in context of a multi-sensor setup.

5. Results

In the following, we outline the conducted experiments and present our results to show the effect of the proposed device localization approach but also the influence of the derived location in context of activity recognition. Subsequently, we evaluate our introduced cross-subjects activity recognition approaches. Due to lack of space, we only present the aggregated results of all subjects. Further, we also removed some tables which belong to the original publication [11]. However, the individual results of each subject and classifier are available[8]. Unless otherwise specified, the provided results are based on the random forest classifier which turned out to consistently perform better than other classification techniques.

5.1. On-body position detection

For the first experiment, we evaluated an activity-independent approach to create a baseline. Thus, we trained for each subject a single classifier on the data of all performed activities and each position. Table 2 shows the result and illustrates that the device position can be recognized with an F-measure of 81%. In this context, the *shin* (P_4) has the highest (88%) and the *forearm* (P_2) and *upper arm* (P_6) the lowest (79%/78%) recognition rate. The latter highlights the problem regarding

[8] http://sensor.informatik.uni-mannheim.de#results2016activity.

Table 2
Activity-independent position recognition rates for different on-body locations.

Class	Precision	Recall	FP-Rate	F-measure
P_1	0.79	0.82	0.04	0.80
P_2	0.79	0.78	0.03	0.79
P_3	0.79	0.82	0.04	0.80
P_4	0.90	0.86	0.02	0.88
P_5	0.83	0.80	0.03	0.82
P_6	0.79	0.78	0.03	0.78
P_7	0.79	0.81	0.04	0.80
avg.	0.81	0.81	0.03	0.81

Table 3
Position recognition rates for static activities and different feature sets: Shows that orientation and time-based features are needed for an accurate recognition.

Features	Precision	Recall	FP-Rate	F-measure
Time-based features	0.72	0.72	0.05	0.72
With orientation	0.88	0.88	0.02	0.88
Only orientation	0.54	0.53	0.08	0.54

Table 4
Detailed results for the proposed position recognition method.

Class	Precision	Recall	FP-Rate	F-measure
P_1	0.87	0.89	0.11	0.88
P_2	0.87	0.85	0.15	0.86
P_3	0.86	0.89	0.11	0.87
P_4	0.95	0.92	0.08	0.94
P_5	0.91	0.90	0.10	0.91
P_6	0.85	0.84	0.16	0.85
P_7	0.91	0.92	0.08	0.92
avg.	0.89	0.89	0.11	0.89

the flexibility of the arm during each activity and also indicates that these two positions are the most problematic device locations. Examining the confusion matrix, shows that the individual positions are not confused. Indeed, the false-positives and the false-negatives are almost evenly distributed.

Further investigations point to the fact that the recognition rate of the correct device location is higher if the related activity is characterized by stronger acceleration. Hence, the separation between static and dynamic activities results in significantly different recognition rates for these two kinds of activity groups (72%/89%).

We examined the feature set and figured out that the gravity vector of the device provides useful information. However, attention should be paid to the fact that our experiments also showed that the gravity vector and derived features (roll and pitch) lead to overfitting. If a classifier was trained for a specific position then the position recognition rate dropped after the device was reattached for this position. This is mainly because the orientation of the device was slightly changed by the subject. The orientation seems not to be a reliable indicator of the current device position. However, investigations have shown that static activities and the device orientation are correlated. Thus, the orientation enables to separate implicit between the static activities which results in less misclassifications of the device position across these activities. In this context, we only considered the introduced discretized orientation. Table 3 summarizes the results and shows that the recognition rate of the device localization in context of static activities increases by 16%.

Certainly, the usage of different feature sets for these two kinds of activity groups (static and dynamic) requires the ability to separate between them. Hence, we constructed a classifier that decides to which activity group the performed activity belongs. The results clearly show that the segmentation performs very well (F-measure: 97%).

As a result, we evaluated the approach where we first decide if a static or dynamic activity is performed and then apply an activity-level specific position classifier. Compared to the baseline, Table 4 shows that this approach has an 8% higher recognition rate. In this context, the *shin* is still the best (94%) and the arm (*forearm* and *upper arm*) the worst (86%/85%) position. Looking at the confusion matrix still exposes an evenly distribution of the false-negatives and false-positives but certainly lower values. This indicates that the distinction of the activity-levels, more precise, the individual handling of the dimensions of the data leads to a better distinction of the device positions. Hence, the experiments show that in most of the cases it is possible to recognize the device position correctly. In general, the considered positions seem not to be confused concerning the classification which confirms that each position provides different information for the same activity.

Table 5
Results of the baseline method for activity recognition without position information.

Class	Precision	Recall	FP-Rate	F-measure
A_1	0.84	0.76	0.02	0.80
A_2	0.77	0.81	0.04	0.79
A_3	0.99	0.94	0.00	0.96
A_4	0.83	0.83	0.03	0.83
A_5	0.68	0.77	0.06	0.72
A_6	0.70	0.64	0.05	0.67
A_7	0.93	0.89	0.01	0.91
A_8	0.85	0.87	0.03	0.86
avg.	0.80	0.80	0.03	0.80

Table 6
Confusion matrix for the baseline activity recognition method without position information.

	Predicted							
	A_1	A_2	A_3	A_4	A_5	A_6	A_7	A_8
A_1	4997	910	2	3	41	23	52	554
A_2	514	6758	1	36	155	108	36	784
A_3	5	2	1114	0	0	0	66	0
A_4	14	94	0	7208	512	837	63	5
A_5	20	108	0	370	6652	1231	224	12
A_6	19	117	0	1000	1798	5622	150	15
A_7	69	95	6	52	611	177	8712	22
A_8	290	741	0	3	49	23	17	7677

In summary, our position-recognition approach that makes use of a random forest classifier and distinguishes between different activity levels achieves an average performance of 89% across all positions.

5.2. Subject-specific activity recognition

The whole idea of our work is based on the idea that knowledge about the device position improves activity recognition. Therefore, we also have to show that the position-aware activity recognition approach that uses the automatically detected device position outperforms the baseline approach that does not consider the device position. For this purpose, we constructed and examined the introduced position-independent activity classifier for each subject which was trained on all data of all positions. Table 5 illustrates the performance of this approach and shows that the correct activity is recognized with an F-measure of 80%. However, considering the individual activities, it shows that the recognition rate is unequally distributed. Thus, *sitting* (A_6) has a significantly worse (67%) and *jumping* (A_3) a much better (96%) recognition rate. Additionally, the activities *climbing down* (A_1) and *standing* (A_5) are often confused with other activities. In this context, the corresponding confusion matrix (see Table 6) emphasizes that the recognized activity is often wrong if a performed activity is similar to another, i.e., *lying* (A_4), *standing* (A_5), and *sitting* (A_6) but also *climbing up* (A_1), *down* (A_2), and *walking* (A_8) are often confused.

In contrast, the introduced position-aware approach achieves a 4% higher F-measure. Table 7 shows that in case of each activity, the consideration of the device localization results in a higher or equal recognition rate. Concerning the static activities, we can observe that the F-measure values increased notably. Indeed, the activities *lying* (+6%), *standing* (+5%), and *sitting* (+9%) have improved the most. The related confusion matrix (see Table 8) makes clear that the problem of misclassification is not completely solved but better handled than before. For dynamic activities, the recognition rate improved slightly.

Considering the activities and positions in detail (see Table 9), it leads to the fact that there is no optimal device position. The chest, waist, thigh, and shin perform on average at best but they perform different depending on the activity. Thus, the activity *climbing stairs up* is best handled by the chest (up to 5% better) whereas the *thigh* recognizes the activity *standing* the best (up to 14% better). This confirms a statement of a previous work where they stated that the optimal sensor placement depends on the activity [20]. Further, it points out that most of the positions perform still bad regarding the static activities. This indicates that even low acceleration combined with the (predicted) device position makes it hard to distinguish between such activities. Besides, there are also activities where each position performs very well. Hence, the activities *running* (≥91%) and *jumping* (≥95%) are equally well recognized for all positions due to the high acceleration of the devices. These show that the acceleration strength is decisive concerning the activity recognition rate and that in case of low acceleration additional information of the environment or context-related information are required.

Despite the fact that we recognized only in 89% of all cases a correct device position and compared with the position-independent approach (80%), these results indicate clearly that the consideration of the device position results in a higher

Table 7
Results of the proposed activity recognition method that uses automatically detected device positions.

Class	Precision	Recall	FP-Rate	F-measure
A_1	0.84	0.77	0.02	0.81
A_2	0.78	0.81	0.04	0.79
A_3	0.99	0.95	0.00	0.97
A_4	0.90	0.88	0.02	0.89
A_5	0.74	0.81	0.05	0.77
A_6	0.78	0.74	0.04	0.76
A_7	0.94	0.91	0.01	0.92
A_8	0.85	0.88	0.03	0.86
avg.	0.84	0.83	0.03	0.84

Table 8
Mean confusion matrix: Proposed activity recognition method using automatically detected device position.

	Predicted							
	A_1	A_2	A_3	A_4	A_5	A_6	A_7	A_8
A_1	5080	849	2	4	42	24	40	548
A_2	526	6820	1	26	134	87	31	768
A_3	7	5	1130	0	0	0	46	1
A_4	18	94	0	7660	324	579	57	8
A_5	19	99	0	217	7000	1020	244	15
A_6	19	112	0	582	1380	6470	141	18
A_7	70	96	11	38	535	142	8830	24
A_8	287	709	1	3	50	24	14	7720

Table 9
Results of the proposed activity recognition method with known device positions.

Class	P_1	P_2	P_3	P_4	P_5	P_6	P_7
A_1	0.86	0.75	0.76	0.83	0.81	0.80	0.82
A_2	0.83	0.72	0.76	0.84	0.83	0.78	0.80
A_3	0.97	0.97	0.97	0.95	0.95	0.98	0.97
A_4	0.89	0.83	0.89	0.90	0.86	0.94	0.91
A_5	0.72	0.73	0.71	0.86	0.84	0.75	0.81
A_6	0.72	0.76	0.65	0.82	0.80	0.74	0.82
A_7	0.92	0.91	0.91	0.93	0.94	0.92	0.93
A_8	0.89	0.82	0.82	0.89	0.88	0.85	0.88
avg.	0.84	0.80	0.79	0.87	0.86	0.83	0.86

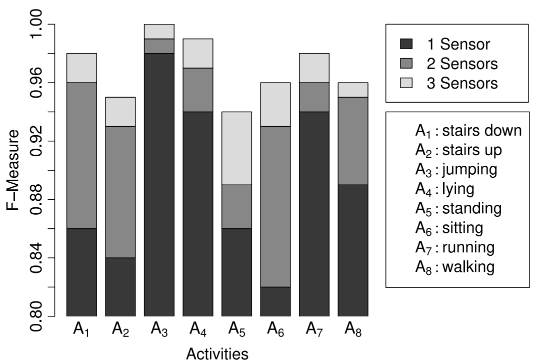

Fig. 5. The recognition rates of a multi-sensor setup. It illustrates the possible improvements of the recognition rate for each activity.

activity recognition rate (84%). So, the results show that it does not depend on the activity but on the device position if the information of the device position improves the activity recognition rate. In this context, also the individual handling of the different dimensions (e.g., device position and activity-level) leads to a better distinction of the target classes, so to a better recognition rate. Especially in the context of the static activities, these two approaches lead to a significant better recognition.

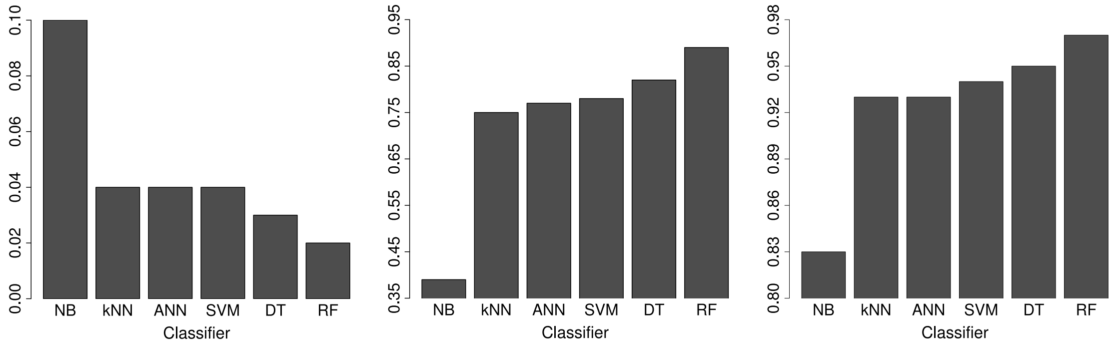

(a) FP-Rate. (b) F-measure. (c) Accuracy.

Fig. 6. Performance of the different classifiers for position recognition in the activity-level dependent scenario.

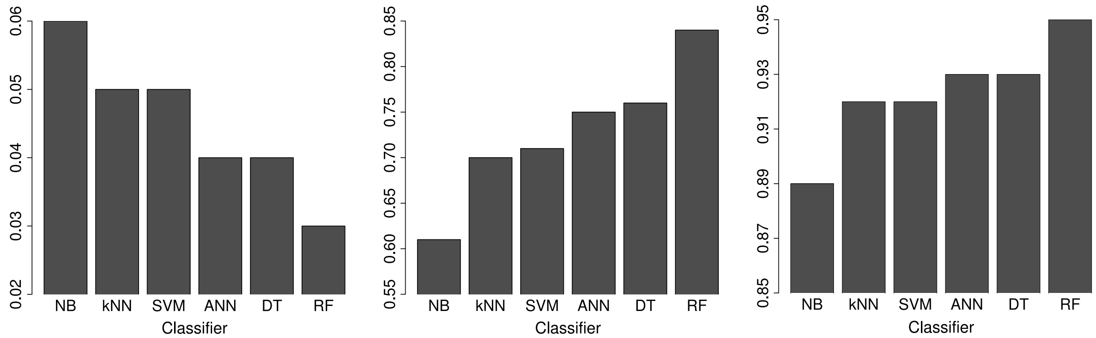

(a) FP-Rate. (b) F-measure. (c) Accuracy.

Fig. 7. Performance of the different classifiers for position-aware activity recognition. The on-body device position was detected in a previous step by the activity-level dependent approach (Random Forest).

If we shift our focus to a scenario where we could rely on additional wearable devices, Fig. 5 shows the improvements concerning the different activities. Indeed, considering all activities, a two-part setup performs always better then a single sensor independent of the selected on-body device positions. Hence, the worst two-part setup (head and upper arm) still achieves a recognition rate of $\geq$90% where the best combination (tight and waist) has up to 94%. Besides, the worst combinations always cover a position which is located on the arm or on the head. This is consistent with the preceding results, i.e., it is due to the flexibility. In contrast, the best two-part combinations consist always of the sensors which performed the best in a single sensor environment. All of this also holds if we compare a three- and two-part setup.

Considering the individual activities, the biggest improvements with a two-part setup could be achieved concerning *sitting* (A_6, +11%), *climbing stairs* (A_1, +10% and A_2, +9%) and *walking* (A_8, +6%). This is strong evidence that already one additional wearable device increases the robustness and quality of the recognition system significantly. Further, it does not matter if the on-body position selection is up to the subject. A third sensor still improves the recognition for all activities but less significant.

5.3. Comparison with other classification methods

In order to show the benefits of using the proposed random forest classifier, we compared its performance to the one of other common classification methods, in particular Artificial Neural Network (ANN), Decision Tree (DT), *k*-Nearest Neighbors (kNN), Naive Bayes (NB), and Support Vector Machine (SVM). All of these classifiers were used in previous works on activity recognition and achieved good results.

Considering the activity-level dependent position recognition approach, the other classifiers performed worse. Fig. 6 illustrates the results and shows clearly that *Random Forest* (89%) outperforms the other classifiers. In this context, *NB* (39%) performed the worst probably due to assumption that all features are independent. In contrast, *k-NN* (75%), *ANN* (77%), and *SVM* (78%) achieved reasonable results. We performed parameter optimization and choose a radial basis function regarding *SVM*. The *DT* (82%) performed second best but the recognition rate is much worse (−7%). Besides, the training phase of the *RF* was one of the fastest whereas *ANN* and *SVM* took the longest.

Concerning activity recognition, we evaluated the performance of the classifiers in context of position-aware activity recognition based on the recognized device positions of the random forest. Fig. 7 shows that *RF* (84%) achieved the highest activity recognition rate where *NB* (61%) performed the worst. Further, *k-NN* (70%) and *SVM* (71%) performed almost

Table 10
Dynamic activity recognition (F-measure): Performance of cross-subjects approaches on each individual device position. Each classifier was only trained and tested with data of a specific on-body position.

Position	Randomly	L1O	Top-Pairs	Physical
P_1	0.64	0.70	0.69	0.68
P_2	0.60	0.66	0.64	0.65
P_3	0.56	0.62	0.61	0.61
P_4	0.63	0.70	0.71	0.70
P_5	0.54	0.58	0.58	0.59
P_6	0.65	0.72	0.71	0.72
P_7	**0.69**	**0.76**	**0.77**	**0.78**

Table 11
Dynamic activity recognition rate (F-measure) for each cross-subjects approach: The classifiers were trained on data that belongs to the waist (P_7).

Class	Randomly	L1O	Top-Pairs	Physical
A_1	0.62	0.65	**0.69**	**0.69**
A_2	0.62	**0.70**	**0.70**	**0.70**
A_3	0.75	**0.83**	0.82	0.78
A_7	0.87	0.89	**0.92**	0.91
A_8	0.63	0.76	0.75	**0.78**
avg.	0.69	0.76	0.77	**0.78**

equal but worse than *ANN* (75%) and *DT* (76%). Besides, we also evaluated the performance of all classifiers in a position-independent scenario but it exposed that independent of the classification technique the position-aware approach is always better.

These results show that the use of the random forest classifier is not only the best classification method for determining the device position, it also outperforms all other classifiers with respect to determining the activity given a hypothesis about the position of the device.

5.4. Cross-subjects activity recognition

In several cases, people are unable to collect and label data which is required for a subject-specific approach. Therefore, we also focused on the feasibility to recognize the performed activity and device position by relying only on labeled sensor data of other people. For that purpose, we evaluate the performance of the introduced cross-subjects approaches *randomly*, *leave-one-subject-out* (L1O), *top-pairs*, and *physical*. We aim to clarify how differently these approaches perform but also the performance in general depending on the device position and compared to a subject-specific approach. During the first experiments, we only consider dynamic activities as target classes to avoid misinterpretation. Thus, we assume that static activities are less characterized by an individual person, i.e., the subtle acceleration that is performed by these activities is probably similar for many different groups of people.

As a first step, we focused on the activity recognition rate of position-dependent classifiers to expose differences in performance. Table 10 shows that across all positions the introduced approaches perform comparable but the recognition rate varies significantly. The waist seems to be the best on-body position for all approaches where *physical* achieves the highest activity recognition rate (78%). In this context, the results indicate that the acceleration patterns for the same activity across several users are most similar at this position. Considering the baseline (*L1O*), *top-pairs* (+1%) and *physical* (+2%) perform slightly better while they have to process significantly less data. Besides, previous work already showed that *L1O* would not scale in a large-user environment due to the varying behavior. Actually, the classifier seems only to learn the dominant behavior across all people, i.e., individual behavior is lost and rate as noise. Considering the other positions, it points out that surprisingly the thigh (P_5) based classifier performs the worst. We examined the individual acceleration patterns and detected that the bad performance results from the unstable position of the device (trouser pocket). Hence, the device was able to move slightly during the data collection. This kind of noise could be handled by a subject-specific approach because it was consistent but this is not the case across subjects. However, this does not mean that the position is unsuitable but, e.g., needs more effort concerning personalization [10].

Considering the recognition rate of the individual activities, Table 11 shows the corresponding recognition rates of the waist-based classifier. Independent of the evaluated approaches, *climbing stairs* (∼70%) has the lowest and *running* (∼91%) the best recognition rate. Indeed, compared to *L1O*, it points out that *physical* recognizes all activities better expect *jumping*. In this context, especially *climbing stairs* and *walking* have a higher recognition rate. This is remarkable because these are the only dynamic activities which are most often confused. We believe that this is evidence for the feasibility to rely on common physical characteristics to identify meaningful groups. However, we also conclude that our considered physical

Table 12
Improvement of the activity recognition rate (physical) with an additional accelerometer. The classifier was trained on data that belongs to the shin (P_4) and waist (P_7).

Class	P_7 (transmitter belt)			P_4–P_7 (add'l smart-band)		
	Precision	Recall	F_1	Precision	Recall	F_1
A_1	0.70	0.67	0.68	0.72	0.74	0.73
A_2	0.71	0.69	0.70	0.72	0.75	0.74
A_3	0.73	0.84	0.78	0.83	0.92	0.87
A_4	0.98	0.92	0.95	0.99	0.92	0.95
A_5	0.69	0.82	0.75	0.74	0.88	0.80
A_6	0.76	0.80	0.78	0.80	0.86	0.83
A_7	0.91	0.78	0.84	0.94	0.79	0.86
A_8	0.77	0.79	0.78	0.83	0.75	0.79
avg.	0.79	0.79	0.79	0.83	0.81	0.82

Table 13
Confusion matrix: Cross-subjects based approach (physical, P_4 and P_7).

	Predicted							
	A_1	A_2	A_3	A_4	A_5	A_6	A_7	A_8
A_1	10 404	2 098	2	1	43	46	383	1 121
A_2	1 561	13 683	0	10	447	456	178	1 828
A_3	13	0	2342	0	0	0	196	0
A_4	6	88	0	17 175	157	1 273	28	5
A_5	7	11	0	1	16 192	1 992	271	1
A_6	2	49	0	138	2 308	16 142	50	3
A_7	780	158	496	41	2 571	380	16 508	35
A_8	1 683	2 847	0	0	100	17	19	14 213

characteristics do not cover the features of *jumping*. Besides, *top-pairs* performs slightly better than *L1O* but, e.g., concerning *walking* even worse. We noticed during the experiments that the acceleration patterns were partly contradictory while the classifier learned the dominant behavior.

Subsequently, we also considered static activities (A_4–A_6). Table 12 shows that the recognition rate seems to be stable but the recognition rate of dynamic activities drops slightly. During this experiment, we also applied the introduced static and dynamic activity split (including all errors) to consider the gravity based feature in context of static activities. On the one hand, this division caused the deterioration of the dynamic activity recognition rate, on the other hand the confusion matrix shows (not presented) that especially lying (A_4) and standing (A_5) are significantly less confused due to the considered gravity based feature. Thus, the results indicate that this feature is also reliable across people. Compared to our subject-specific approach (see Table 9), especially the recognition of *climbing stairs* performs worse whereas the recognition rate of static activities is comparable (±2%). This confirms our initial assumption concerning static activities in context of cross-subjects models.

To address the difference in performance, we also analyzed the improvement that could be achieved by an additional acceleration sensor. After all, several people already wear two devices. Table 12 illustrates the possible improvement if we combine two of the best performing on-body device positions. In average, the recognition rate increases by 3% where especially the recognition of climbing stairs improved (+5%). On the downside, *walking* only increased slightly. However, this also makes clear that this activity is the most challenging while it is the one that is most often performed during the day. In this context, Table 13 shows the corresponding confusion matrix. It strikes that the problematic groups are still *climbing up* (A_1), *climbing down* (A_2), *walking* (A_8) and *lying* (A_4), *sitting* (A_5), *standing* (A_6). Compared to our subject-specific approach, it points out that no new issues arise but existing will become more manifest, e.g., jumping is more often confused with running.

Finally, we also investigated if cross-subjects based models are able to recognize the on-body device position. Table 14 shows the individual recognition rate. Independent of the approach, it points out that the recognition quality differs significantly across the different positions where waist (78%) and shin (74%) are best recognized. Considering the overall results, we have to state that the position recognition rates are not sufficient to be considered as part of an activity recognition system. However, these results also confirm our assumption that the waist seems to be the best on-body device position for cross-subjects activity recognition.

In general, the results show that cross-subjects models are feasible for activity recognition if the on-body device position is known a-priori. In this context, the waist is the best device position for cross-subjects activity recognition where we were able to achieve a recognition rate of 79%. Considering an additional wearable device, improved the performance by +3%. Thus, our results indicate that it is feasible to monitor the physical activities of people which are unable to collect and label

Table 14
Activity-independent position recognition (F-measure): Performance of cross-subjects approaches concerning the recognition of the on-body device position.

Class	Randomly	L1O	Top-Pairs	Physical
P_1	0.56	0.63	0.59	0.61
P_2	0.58	0.63	0.59	0.58
P_3	0.54	0.61	0.56	0.57
P_4	0.68	0.74	0.72	0.73
P_5	0.53	0.60	0.57	0.58
P_6	0.50	0.57	0.53	0.54
P_7	0.74	0.78	0.76	0.77

required data. Further, the *physical* based approach performed the best in context of the most reliable device position where especially walking and climbing stairs are better handled. Besides, we consider the recognition of the device position in a cross-subjects scenario still as open issue which needs further investigations.

6. Conclusion and future work

In this paper, we targeted open issues which go along with activity recognition in a real world scenario. Commonly it is up to the user on which on-body position the device is worn, further, several people, e.g. elders, are unable to collect and label the required amount of data to create a classification model. For that purpose, on one hand, we investigated the feasibility to recognize the wearable device position by relying on a common acceleration sensor and if the recognized position influences the activity recognition rate. On the other hand, we examined the performance of different cross-subjects approaches dependent on the different device positions to clarify their feasibility.

The experiments were performed by relying on our large real world data set that comprises 8 different physical activities of 15 people where for each activity 7 on-body positions were recorded simultaneously. Considering this data, our results show that the best recognition rates were achieved with the machine learning technique random forest. In this context, we were able to recognize in a subject-specific scenario, the on-body device position with an F-measure of 89% where this information, including all errors, improved the activity recognition performance by +4% (84%). Hence, the results provide evidence for the improvement of the activity recognition rate in case that the on-body position is known.

Subsequently, we evaluated several cross-subjects activity recognition approaches where the classifier was trained on labeled data of certain people expect data of the target user. The results show that the waist is the best on-body position for cross-subjects activity recognition due to the fact that the acceleration patterns for the same activity across different users are most similar at this position. Further, the results of the individual approaches show that abstract physical characteristics of subjects enable us to build meaningful cross-subjects classifiers. Considering the most reliable device position, the physical based approach was able to achieve a recognition rate of 79%. Considering an additional wearable device, the recognition rate improves by +3% (82%).

In context of previous works concerning on-body detection, researchers presented lower or equivalent recognition rates and considered less positions and activities. Coskun et al. considered the hand, trousers (thigh), and backpack and achieved a recognition rate of 85% [13]. Furthermore, Vahdatpour et al. considered the same on-body positions as we did expect the chest and considered only the activity *walking* but achieved an accuracy of 89% [5]. This indicates that the consideration of more positions and activities lead to a lower recognition rate as we can see in the results of our first experiments that did not distinguish between static and dynamic activities. However, as presented, due to the individual handling of different activity-level groups, our approach performs significantly better in a real world scenario where people change the orientation, device position, and activity all the time.

In case of cross-subjects activity recognition, most existing works focused on leave-one-subject-out where the opinions tend to state that this approach is not reliable. In this context, Vo et al. [10] clarifies that an increasing number of considered subjects goes along with a decreasing activity recognition performance. We attribute this behavior to the fact that the classifier learns only the most dominant behaviors across people. To counteract this behavior, researchers suggest to rely on specific groups where Lara et al. [14] and Weiss et al. [16] hypothesized that physical characteristics such as gender, weight, and fitness level could be reliable indicators to form groups. In our work, we investigated this hypothesis and our results provide evidence for the correctness where especially the problematic group climbing stairs up, down and walking was better handled. However, we also have to state that the considered physical characteristics did not cover the features of the activity jumping.

As future work, we plan to investigate two aspects. On the one hand, we want to focus on personalization of a cross-subjects model and the associated effort for the target user by using active machine learning. On the other hand, we also want to expand our experiments by considering a larger group of people to identify, e.g., less obvious physical characteristics to improve the performance.

References

[1] H. Martín, A.M. Bernardos, J. Iglesias, J.R. Casar, Activity logging using lightweight classification techniques in mobile devices, Pers. Ubiquitous Comput. 17 (4) (2013) 675–695.
[2] L. Barkhuus, V.E. Polichar, Empowerment through seamfulness: smart phones in everyday life, Pers. Ubiquitous Comput. 15 (6) (2011) 629–639.
[3] L. Sun, D. Zhang, B. Li, B. Guo, S. Li, Activity recognition on an accelerometer embedded mobile phone with varying positions and orientations, in: Ubiquitous Intelligence and Comp, Springer, 2010, pp. 548–562.
[4] L. Bao, S.S. Intille, Activity recognition from user-annotated acceleration data, in: Pervasive Computing, Springer, 2004, pp. 1–17.
[5] A. Vahdatpour, N. Amini, M. Sarrafzadeh, On-body device localization for health and medical monitoring applications, in: International Conference on Pervasive Computing and Communications, IEEE, 2011, pp. 37–44.
[6] A. Bulling, U. Blanke, B. Schiele, A tutorial on human activity recognition using body-worn inertial sensors, ACM Comput. Surv. 46 (3) (2014) 33.
[7] E.M. Tapia, S.S. Intille, W. Haskell, K. Larson, J. Wright, A. King, R. Friedman, Real-time recognition of physical activities and their intensities using wireless accelerometers and a heart rate monitor, in: 11th IEEE International Symposium on Wearable Computers, IEEE, 2007, pp. 37–40.
[8] D. Guan, W. Yuan, Y.K. Lee, A. Gavrilov, S. Lee, Activity recognition based on semi-supervised learning, in: 13th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications, (RTCSA 2007), IEEE Computer Society, 2007, pp. 469–475.
[9] A. Reiss, D. Stricker, Personalized mobile physical activity recognition, in: Proceedings of the 2013 International Symposium on Wearable Computers, ACM, 2013, pp. 25–28.
[10] Q.V. Vo, M.T. Hoang, D. Choi, Personalization in mobile activity recognition system using-medoids clustering algorithm, Int. J. Distrib. Sens. Netw. 2013 (2013) 12.
[11] T. Sztyler, H. Stuckenschmidt, On-body localization of wearable devices: An investigation of position-aware activity recognition, in: 2016 IEEE International Conference on Pervasive Computing and Communications, (PerCom), IEEE, 2016, pp. 1–9.
[12] L. Chen, J. Hoey, C.D. Nugent, D.J. Cook, Z. Yu, Sensor-based activity recognition, IEEE Trans. Syst. Man Cybern. Part C Appl. Rev. 42 (6) (2012) 790–808.
[13] D. Coskun, O.D. Incel, A. Ozgovde, Phone position/placement detection using accelerometer: Impact on activity recognition, in: 2015 IEEE Tenth International Conference on Intelligent Sensors, Sensor Networks and Information Processing, (ISSNIP), IEEE, 2015, pp. 1–6.
[14] O.D. Lara, M.A. Labrador, A survey on human activity recognition using wearable sensors, IEEE Commun. Surv. Tutor. 15 (3) (2013) 1192–1209.
[15] K. Altun, B. Barshan, Human activity recognition using inertial/magnetic sensor units, in: International Workshop on Human Behavior Understanding, Springer, Berlin, Heidelberg, 2010, pp. 38–51.
[16] G.M. Weiss, J.W. Lockhart, The impact of personalization on smartphone-based activity recognition, in: AAAI Workshop on Activity Context Representation: Techniques and Languages, 2012, pp. 98–104.
[17] N. Ravi, N. Dandekar, P. Mysore, M.L. Littman, Activity recognition from accelerometer data, in: Proceedings of the 17th Conference on Innovative Applications of Artificial Intelligence - Vol. 3, IAAI'05, AAAI Press, 2005, pp. 1541–1546.
[18] Z. Wang, M. Jiang, Y. Hu, H. Li, An incremental learning method based on probabilistic neural networks and adjustable fuzzy clustering for human activity recognition by using wearable sensors, IEEE Trans. Inf. Technol. Biomed. 16 (4) (2012) 691–699.
[19] A. Henpraserttae, S. Thiemjarus, S. Marukatat, Accurate activity recognition using a mobile phone regardless of device orientation and location, in: International Conference on Body Sensor Networks, IEEE, 2011, pp. 41–46.
[20] L. Atallah, B. Lo, R. King, G.-Z. Yang, Sensor positioning for activity recognition using wearable accelerometers, IEEE Trans. Biomed. Circuits Syst. 5 (4) (2011) 320–329.
[21] K. Kunze, P. Lukowicz, H. Junker, G. Tröster, Where am I: Recognizing on-body positions of wearable sensors, in: Location-and Context-Awareness, Springer, 2005, pp. 264–275.
[22] J.R. Kwapisz, G.M. Weiss, S. Moore, Activity recognition using cell phone accelerometers, SIGKDD Explor. Newslett. 12 (2) (2011) 74–82.
[23] Y. Liang, X. Zhou, Z. Yu, B. Guo, Energy-efficient motion related activity recognition on mobile devices for pervasive healthcare, Mob. Netw. Appl. 19 (3) (2014) 303–317.
[24] S.J. Preece, J.Y. Goulermas, L.P. Kenney, D. Howard, A comparison of feature extraction methods for the classification of dynamic activities from accelerometer data, IEEE Trans. Biomed. Eng. 56 (3) (2009) 871–879.
[25] L. Breiman, Random forests, Mach. Learn. 45 (1) (2001) 5–32.
[26] T.K. Ho, The random subspace method for constructing decision forests, IEEE Trans. Pattern Anal. Mach. Intell. 20 (8) (1998) 832–844.
[27] J.-Y. Yang, J.-S. Wang, Y.-P. Chen, Using acceleration measurements for activity recognition: An effective learning algorithm for constructing neural classifiers, Pattern Recognit. Lett. 29 (16) (2008) 2213–2220.