

Adaptive Activity and Context Recognition Using Multimodal Sensors in Smart Devices

Sébastien Faye^(✉), Raphael Frank, and Thomas Engel

Interdisciplinary Centre for Security, Reliability and Trust,
University of Luxembourg,
4 rue Alphonse Weicker, 2721 Luxembourg, Luxembourg
{sebastien.faye,raphael.frank,thomas.engel}@uni.lu

Abstract. The continuous development of new technologies has led to the creation of a wide range of personal devices embedded with an ever increasing number of miniature sensors. With accelerometers and technologies such as Bluetooth and Wi-Fi, today’s smartphones have the potential to monitor and record a complete history of their owners’ movements as well as the context in which they occur. In this article, we focus on four complementary aspects related to the understanding of human behaviour. First, the use of smartwatches in combination with smartphones in order to detect different activities and associated physiological patterns. Next, the use of a scalable and energy-efficient data structure that can represent the detected signal shapes. Then, the use of a supervised classifier (i.e. Support Vector Machine) in parallel with a quantitative survey involving a dozen participants to achieve a deeper understanding of the influence of each collected metric and its use in detecting user activities and contexts. Finally, the use of novel representations to visualize the activities and social interactions of all the users, allowing the creation of quick and easy-to-understand comparisons. The tools used in this article are freely available online under a MIT licence.

Keywords: Sensing system · Wearable computing · Activity detection

1 Introduction

In recent years, the growing availability and falling cost of smart devices have opened up a world of opportunities for new applications. Apart from smartphones, these connected objects include a wide range of ultra-portable devices that constantly interact with the users and their environment. Among those *wearables*, the vast majority are smartwatches and activity trackers. These have become very diverse and are equipped with high-performance sensors that allow users to monitor their physical activity in a way never possible before. Their sensors can read metrics from arm or hand movements with an accuracy comparable to specialized experimental devices [1]. These devices include physical sensors that are permanently in contact with the user’s wrist, such as motion detectors (e.g. accelerometers) and environmental monitoring sensors (e.g. light sensors, microphone). Their

ability to monitor other physiological metrics, such as heart rate, leads to new areas of research. Further, the recent arrival on the market of major players, like Apple, Google and Microsoft, has facilitated the development and widespread adoption of sensing applications, opening the way to many new areas, including health, sport, and personal monitoring. According to ABI Research, it is estimated that the global market for wearables will reach 170 million in 2017 [2].

At present, whether we are talking about smartphones or wearables, these connected objects are generally used individually and for specific consumer applications (e.g. fitness). In most cases, the classic data fusion from sensors is adapted to be made in real time (e.g. pattern finding). This requires heavy-duty processing algorithms and consumes energy. Moreover, most systems only use smartphones, whereas wearables are more suitable for detecting user activities. Finally, few studies have looked at all types of existing sensors with the intention of arriving at a scalable and easy-to-implement solution.

In this paper, we intend to go one step further by presenting a sensing system that combines the data collected by one smartwatch and one smartphone. The platform that we have developed makes use of commercially-available devices and can be used to analyse the activity of a monitored user in great detail. Possible applications range from sports tracking systems to human behaviour modeling. Our contribution addresses four complementary objectives. (1) The design of an energy-efficient sensing system, using a streamlined fusion of data collected on two devices (a smartphone and a smartwatch). (2) The use of a supervised machine learning model to recognize user activities and their contexts. (3) The combination of multimodal metrics to obtain more advanced feature sets and generalize the recognition process. Finally, (4) the comparison of activities and social interactions of different users using novel 3D visual representations.

In the following section we provide a review of existing literature. Next, in Sect. 3, we present of our sensing system, which is focused on the devices used for data collection and on how they communicate to exchange data. Section 4 describes our experimental campaign and how we used the collected metrics to create the data set used for our analysis. Section 5 focuses on the analysis of the data set and presents some relationships between metrics and a set of predetermined activities using a Support Vector Machine (SVM) model. These relationships form the basis for the recognition of activities and contexts to be inferred. Finally, two profile comparison methods are introduced in Sect. 6, before we conclude in Sect. 7.

2 Related Work

The use of mobile devices as key elements in a sensing system has been discussed for many years, both in industrial and research communities, as an opportunistic [3] or a participative system [4]. The classic architecture for such a sensing system consists of three parts [5,6]. First, individual devices collect the sensor data. Then, information is extracted from the sensor data by applying learning methods, generally on one of the devices or in the cloud, depending on the sensitivity of the data, the sampling strategy or the privacy level applied. Finally, the data can be shared and visualized from the cloud.

Smartwatches have their place in this kind of architecture and can open up new perspectives as they can collect the user's activity and physiological signals [7], while smartphones are reserved for recording the user's context. Smartwatches and smartphones are usually connected via Bluetooth Low Energy [8], a relatively new technology standardized under the Bluetooth 4.0 specification [9]. Compared to smartwatches, smartphones have a better battery capacity and can launch several tasks at the same time. By using a smartphone as a local gateway to access the Internet – via Wi-Fi or Cellular – we can turn this local sensing platform into a connected ecosystem [6].

As the applications need to be running on the devices permanently to collect and send data, there is an important compromise to be found between sample rate, rate of transmission and the consumption of energy [8]. The authors of [10] show, for example, that using all the sensors of a LG Nexus 4 E960 can reduce its battery from 214.3 h (no sensors) to 10.6 h (all sensors). Some systems attempt to circumvent this energy limit by offloading data processing onto servers [11]. Others propose sharing the data among neighboring phones [12]. By these means, cloud computing is widely used with smartphones and allows the creation of elastic models [13], where applications are launched on the mobile phone, and the data is processed in the cloud.

In the surveyed literature, accelerometers are the sensors most commonly used to recognize various physical and upper body activities. Indeed, [1] shows that specific movements of the arms, the hands and the fingers, generate sufficient energy to be distinguished by the accelerometer and the gyroscope in a smartwatch with 98% precision. By correlating different sources of data, other sensors such as GPS, microphones and Wi-Fi signals can also be used to improve the classification accuracy and estimate, for example, the mode of transport (e.g. bike, car) [14]. By continuously recording sound, it is possible to identify different user contexts, whether having a conversation, sitting in an office, walking out on the street or even making coffee [15, 16]. SPARK [17] is a framework that can detect symptoms associated with Parkinson's disease using a smartwatch on the wrist (to detect dyskinesia using motion sensors), and a smartphone in the pocket (gait analysis and sound). Shin et al. [18] study patients with mental disorders and use smartwatches to help quantify the exercise and the amount of sunlight wearers have received, using GPS, accelerometer and the light sensor. Video sensing also permits various activities to be recognized [19]. However, video analysis is both algorithmically and computationally expensive, especially in a resource-constrained environment. Finally, social interactions can be identified using Bluetooth, Wi-Fi, Near-Field Communications (NFC) or cellular [10].

Activity detection involves the recognition of spatio-temporal patterns from sensor data that is usually incomplete and noisy. There is significant number of models that are able to characterize human behaviour from different features (e.g. accelerometer data). The temporal signal shape can be analyzed both in time and frequency domains. Time-domain features include basic waveform characteristics and signal statistics that can be considered to be features of a given signal, e.g. the statistical moments, time between peaks, binned distribution,

Table 1. Specification of the devices used of our studies.

Devices	RAM / Storage	CPU	Battery	Network Interfaces	Collectable Data
Samsung Gear Live	512 MB / 4 GB	Quad-core 1.2 GHz	300 mAh	Bluetooth 4.0	Heart rate, heart rate accuracy, pedometer, linear acceleration, time.
LG Nexus 5	2 GB / 16 GB	Quad-core 2.3 GHz	2,300 mAh	4G/LTE, GPS, 802.11 a/b/g/n/ac dual-band Wi-Fi, NFC, Bluetooth 4.0	Proximity, ambient light, linear acceleration, ambient sound, detected activity, activity confidence, pedometer, mobile network information, detected Bluetooth devices, Wi-Fi networks, GPS (location, altitude, speed, bearing, distance), local weather (from OpenWeatherMap.org), time.

mean value of local maxima [20]. Data set reduction techniques such as Principal Component Analysis and Linear Discriminant Analysis can be used to extract the most significant discriminating features while reducing the dimensionality of the data representation [21]. Combining the feature extraction techniques above, activity recognition can be trained using (semi-)supervised methods in a controlled setting. These methods include Decision Trees, Neural Networks and Support Vector Machines, all of which have been successfully used in human activity recognition [22]. For example, Frame-based Descriptor and multi-class Support Vector Machine [23] is an approach that can classify a large variety of gestures. Unsupervised methods (e.g. k-means clustering [24]) can then be used to find structures in the different activity sequences and durations that were identified to find common properties or behaviours of user groups.

3 Sensing System

In order to carry out our studies and obtain the results presented in this article, we used our own system, SWIPE [7], which is available online¹ under a MIT licence. It is composed of two main parts: an Android application for data collection and a web platform for data processing.

3.1 Hardware

We used two devices running Android 5.1.1. One was a smartwatch (Samsung Galaxy Gear Live) that records the wearer’s activity by registering wrist movements and other physiological data (i.e. heart rate). The other, a smartphone (LG Nexus 5), is responsible for collecting contextual data (e.g. with its microphone) as well as some additional activity data (e.g. accelerometer). The decision to run SWIPE on Android makes sense because of its maturity and its leading role in the current generation of smartwatches. Table 1 summarizes the specifications of the two devices, including details of the data that our system is able to collect.

¹ <https://github.com/sfaye/SWIPE/>.

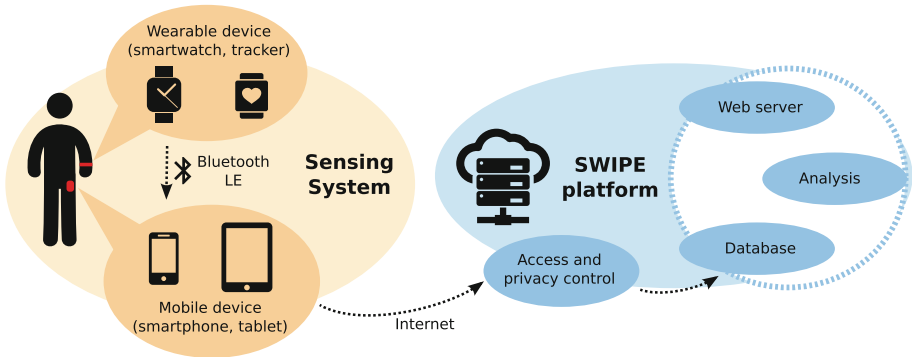


Fig. 1. SWIPE overall architecture.

3.2 Architecture

The architecture of SWIPE is shown in Fig. 1 and consists of two parts. First, the sensing system is composed of a watch (worn on the wrist) and a phone (carried in the pocket) as introduced in Sect. 3.1. The watch periodically sends the data it has collected to the smartphone, which acts as a local collection point and as a gateway to access the SWIPE platform over the Internet (via Wi-Fi or a cellular network). The SWIPE platform is divided into several modules, which (1) receive data following authentication and (2) store, (3) analyse and (4) display the data by means of a web interface. Each user is identified by a unique hash string and his or her data is stored on an internal University of Luxembourg server, which is accessible only on the local network. The link between the server and the sensing system is performed by an intermediate server that acts as a relay.

3.3 Metrics Collected by SWIPE

The main metrics that our system collects are shown in Table 2. The “recording rate” column indicates the frequency at which a metric is saved in a data file, while the “sampling rate” indicates the frequency at which the system acquires raw data from sensors. Since the user is wearing the watch all the time, metrics associated with the watch include the ability to recognize activity. The average speed of movement of the user’s arm is recorded every 30 seconds, along with the maximum speed in order to detect sudden, unusual gestures. Metrics collected by the phone include contextual data. This includes accelerometer readings that are complementary to those provided by the watch. We also store microphone readings to register the level of ambient noise, enabling us to distinguish between noisy and a quiet places. Network data also enables us to collect information on both mobility (GPS, Wi-Fi) and interaction with other users (Bluetooth).

Table 2. Key metrics collected by SWIPE.

Devices	Metrics	Sensors	Recording & Sampling rates		
Phone, watch	Maximum acceleration	Accelerometer	30 sec.	< 1 sec.	Maximum value of $\alpha = \sqrt{(x^2 + y^2 + z^2)}$ $m.s^{-2}$, where x, y and z are the acceleration along each axis of the device, excluding gravity.
	Average acceleration	Accelerometer	30 sec.	< 1 sec.	Average value of α .
	Pedometer	Accelerometer, Android API	60 sec.	~	Number of steps taken by the user, detected by the Android system as a function of the accelerometer.
Watch	Heart rate	Optical heart rate sensor	60 to 300 sec.		Heart rate, in beats per minute, provided by the optical heart rate sensor.
Phone	Ambient sound	Microphone	60 sec.	1 sec.	Ambient sound level, from [0 : 100].
	Bluetooth devices	Network	120 sec.		Number of Bluetooth devices.
	Wi-Fi APs	Network	300 sec.		Number of Wi-Fi Access Points.
	Mobile network data state	Network	300 sec.		Value expressing the use of cell phone network. A reading of zero indicates that the phone is connected to a Wi-Fi AP. For our experiments, only the access point of our workplace was configured.
	Speed	GPS	60 sec.		Travel speed, in $km.h^{-1}$.

3.4 Energy Saving Strategy

The provision of a sensing system launched as a background service represents a potential burden on the batteries of the devices used, which (particularly in the case of smartwatches) are not renowned for their longevity. It is therefore critical that we make every effort to save energy. This includes finding the right compromise between energy consumption and data collection. The proposed system aims to run uninterrupted for at least 12 h in order to collect enough data to obtain an overview of daily activities. To achieve this, we implemented the following optimization strategy.

(1) Data transmission consumes a significant amount of energy. We first configure our application so that the watch, if close enough, uploads its data to the smartphone every 20 min rather than continuously. This allows the application to automatically turn off Bluetooth most of the time and makes the watch fully autonomous (i.e. the user can wear the watch without having to carry the phone). Data collected and transmitted by the smartwatch is received and stored locally by the smartphone to be sent once a day to our servers for later analysis. The data is sent at a predefined time (at midnight) or when the battery level of either of the devices drops below a threshold of 5%.

(2) Another factor that contributes to energy consumption is the frequency at which the sensors record data. The higher the frequency and the longer the transmission time, the more energy is consumed. On the other hand, a lower data acquisition rate will dilute the quality of the resulting data set. Consequently, each metric is configured with the parameters set out in Sect. 3.3. Note that while most of the metrics are configured with a fixed and adequate sampling frequency with respect to the tests carried out, other strategies are set up for specific cases.

Indeed, the acquisition frequency of the heart rate sensor is designed to adapt to the activity of the user. When the user is making little or no movement, the sampling frequency is low, since his heart rate should be stable and the measurements reliable. Conversely, when the user moves, the sensor becomes more sensitive and his heart rate is likely to change. In this case, the data acquisition rate increases in order to take more probes.

(3) Finally, the devices are configured to prevent users from interacting with them. Each is locked with a password and all the unnecessary services managed by Android, such as notifications, are disabled. This allows us to record the data without interruption and under the same conditions for every participant.

This energy saving strategy is evaluated by comparing it with the settings where transmission, harvesting and recording frequencies were high (i.e. all set to 1 second). We find an autonomy gain of about 287% for the smartwatch (13.5 h vs. 4.7 h) and on the order of 189% for the smartphone (15.7 h vs. 8.3 h).

4 Building a Data Set

4.1 Scenario

The studies we conducted involved 13 participants working in the same building at the University of Luxembourg. These participants were selected as a representative sample of both genders and of different ages. Each participant was systematically subjected to the same requirements: (1) wear the watch and smartphone for one day, from 10:00 to 23:59; (2) complete a questionnaire² asking for an exact description of activities carried out (work, commute and leisure activities); (3) sign an informed consent form to accept the privacy policy of the study.

4.2 Example

Figure 2 shows data from one of the participants over a period of about 14 h. The accelerometer data and the level of ambient noise immediately reveal several distinct situations. Around 19:00, for example, the participant appears to perform much faster movements than usual with both his watch and his phone – indicating that he is carrying both devices. The noise level is also high, indicating either a noisy place or excessive friction (which is common when the phone is carried in a pocket). We can easily deduce that the user was running. This is confirmed by the activity recognition algorithm provided by Android, which is able to detect basic activities. The situation is similar around 18:00. The environmental noise level is high, but both devices detect much less movement and the GPS records more rapid progress from place to place: the user was driving. These initial observations form the basis of our intuitive understanding of the user’s activity.

² Available online: <http://swipe.sfaye.com/mobicase15/questionnaire.pdf>.

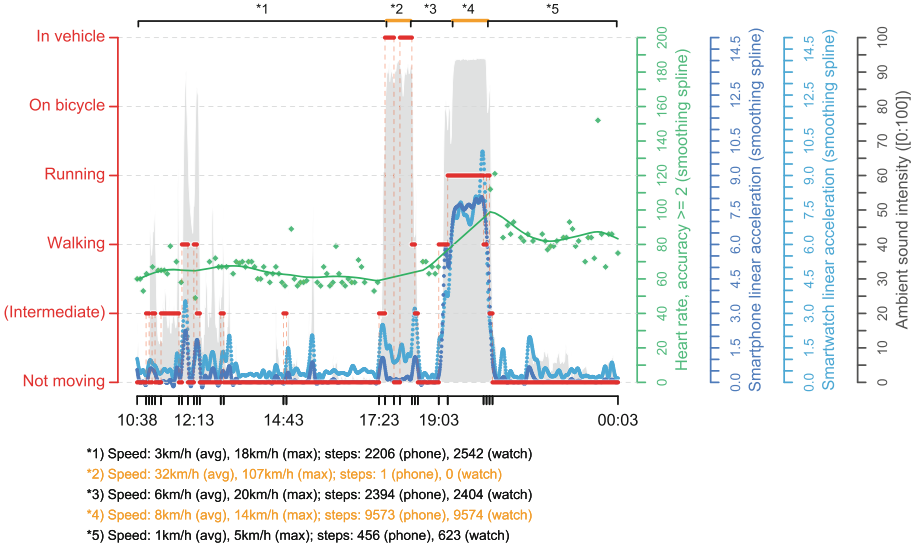


Fig. 2. Example of collected metrics for one participant.

Table 3. Identified activity and context classes with their total durations in our data set, which consists of 157.2h of recordings.

Activity									Context					Total
Physical					In vehicle				Working	Meeting	Shopping	Break	At home	
Sitting	Standing	Walking	Running	Tennis	Bus	Train	Motorcycle	Car	Working	Meeting	Shopping	Break	At home	65.4h
5.2h	11h	3.8h	1.9h	0.8h	1.4h	0.6h	0.3h	7.8h	5.2h	8.1h	3.4h	8.7h	7.3h	

4.3 Activity and Context Classes

In order to build a data set, we used both the information provided by users in the questionnaire and the information from the sensing platform. Each participant told us about the activities he or she had performed. By gathering all the information from the 13 participants, we obtained a total of nine activities (i.e. sitting, standing, walking, running, playing tennis, on a train, on a bus, on a motorcycle, in a car) that can be classified within five different contexts (i.e. working in an office, attending a meeting, in a shopping centre, on a break at work, at home), as represented in Table 3. Since we have the time slots for each activity (e.g. Figure 2), we are able to assign a set of representative values for each activity and context considering multiple inputs.

5 Activity and Context Recognition Using SVM

5.1 Parameters

The problem to be solved is how to identify a class based on a set of metrics. We chose to use SVM (Support Vector Machine) [25], a set of supervised learning techniques to classify data into separate categories. SVMs have the ability to deal with large amounts of data while providing effective results. They can be used to solve problems of discrimination, i.e. deciding which class a sample is in, or regression, i.e. predicting the numerical value of a variable.

For our evaluation we used the SVM classifier provided by the `e1071` package for R [26]. The default optimisation method – C-classification – is used, as well as the classic radial kernel. Grid-search with 10-fold cross validation [27] was used to adjust the cost parameter C (within a range of 1 to 100), as well as γ (within a range of 0.00001 to 0.1).

5.2 Feature Set

The numerous measurements that we have in our data set were not all recorded at the same frequency. As shown in Table 2, acceleration was recorded twice as often as GPS speed. To simplify future operations, we chose to refine the data for each metric by sampling the same number of values from each. For each of the known classes selected in Sect. 4.3, we use a sliding window of ten minutes, moving over the data stream every five minutes. With each movement of the window, two representative values of data included in the window – referred to as x – are recorded: their average \bar{x} , which gives an overall view of the data over the interval; and their standard deviation $\sigma(x)$, which is fundamental to understanding the variations around the average. Finally, each activity and context class is represented as a set M of m metrics, each of which is represented, for each 10-minute data interval x , as \bar{x} and $\sigma(x)$. The following matrix illustrates the structure of the data set:

$$T_1 \begin{pmatrix} \textit{Class} & M_1^{avg} & M_1^{std} & \dots & M_m^{avg} & M_m^{std} \\ \textit{class}_1 & \bar{x}_{1,1} & \sigma(x_{1,1}) & \dots & \bar{x}_{1,m} & \sigma(x_{1,m}) \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ T_n & \textit{class}_\Delta & \bar{x}_{n,1} & \sigma(x_{n,1}) & \dots & \bar{x}_{n,m} & \sigma(x_{n,m}) \end{pmatrix} \quad (1)$$

This representation is simple and has the advantage of abstracting from the excessive precision of the data. It also has the advantage of being lighter and less expensive to treat with a classification algorithm. Assuming we have a set of data composed of t seconds of recording, that the length of the sliding window is t_{window} seconds and that it moves every $t_{step} \leq t_{window}$ seconds, we obtain a data matrix whose size is:

$$columns = (2 \cdot m + 1) \quad rows = \frac{t - (t_{window} - t_{step})}{t_{step}} \quad (2)$$

Our activities database contains, for example, a total of 65.4h of recordings and is 19×784 in size.

5.3 Recognition Using Metrics Individually

First of all, we investigate the individual influence that each metric can have on the recognition of an activity and/or context. Figure 3 represents some selected normalized metric averages over all participants and for each class. For reasons of visualization, the vehicle activities are grouped into the “In vehicle” class. The colour transition between each class represents half the distance that separates their average. The findings are logical, but they confirm the individual importance of each metric. For example, on average the GPS speed reading can help to detect whether the user is traveling in a vehicle, running or at rest. Maximum accelerometer readings can help us recognize a sport activity, such as tennis. Noise in a shopping centre seems to be higher than noise during a meeting.

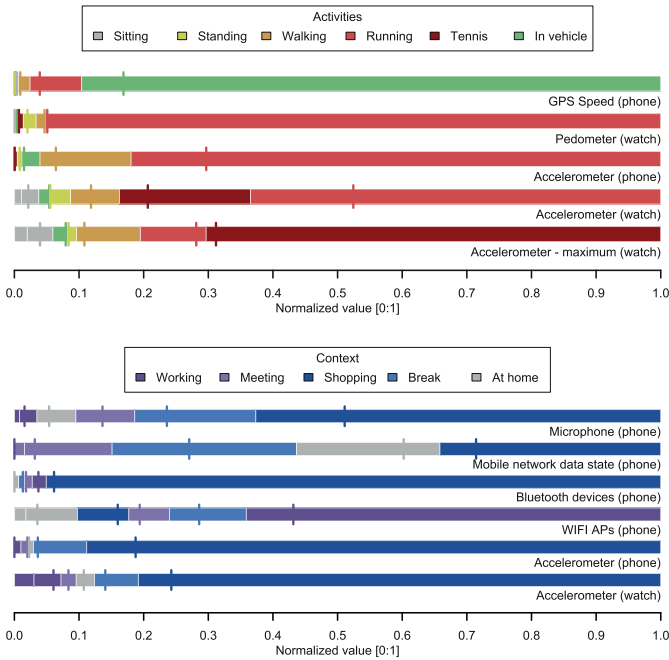


Fig. 3. Selected metric averages for each class.

We want to use streamlined versions of the data set described in Sect. 5.2, with the aim of representing each metric individually to see whether or not it can accurately detect a class. Each data set is evaluated in order to discover how accurately we can predict a class based on a single metric. To do this, each data set is randomly divided into two parts. The first is the training set, comprising 70 % of instances. The second is the test set, comprising the remaining 30 %. The training set is subjected to a grid search to find the cost and γ that minimize the error rate. An SVM model is created from the training set using the best

cost and the best γ . The model is then confronted with the test set with the aim of predicting the number of instances in the test set whose class is correctly recognized by the training set. In order to ensure a representative average value for the error rate, this operation is performed 100 times for each combination and calculated for each iteration as $1 - Accuracy$. The results are shown in Table 4 with $Accuracy = \frac{\text{true value}}{\text{total value}}$.

We notice a huge disparity between all combinations of metrics and classes. The overall findings were quite polarized: some metrics can identify a class with very high reliability (e.g. the relationship between acceleration and running), while others cannot. Of course, the combinations shown are representative of our data set, where activities were taking place in an urban environment. For example, it is normal to see Wi-Fi sometimes taking particular prominence. This would probably not be the case if environments were more heterogeneous.

5.4 Recognition Using a Combination of Multiple Metrics

It makes sense to use a classifier such as an SVM when combining multiple metrics to deduce an activity, which can be seen as a more advanced feature set. We are interested in minimizing the error rate returned by an SVM model, that takes a set of metrics as its input, i.e. finding a combination that minimizes the error rate for both of the activity category and the context category. To do this, we generate all possible combinations of metrics and create a data set for each combination (e.g. watch acceleration and heart rate, Wi-Fi access points and GPS speed and Bluetooth devices, etc.). In the same manner and with the same parameters as above, for all possible combinations, each data set is randomly divided into a test set and a training set in order to calculate the average error rate provided by the combination, over 100 iterations. The combination retained is the one with the minimum average error rate.

Table 5 represents the best combination of metrics obtained for each class of activity and context and for three cases: combined watch and phone metrics, watch metrics, and phone metrics. For each line, the best combination presented is the one that has the best accuracy. For example, the best combination for recognizing the “standing” class is a combination of metrics on the watch and on the smartphone, giving us a 95.3% average recognition accuracy. We can also see that, for the “running” and “motorcycle” classes, using the watch alone provides better accuracy than a combination of the watch and phone sensors. However, in most cases, the combined use of both devices offers better results than a phone or a watch alone. On the whole, the conclusions on the dataset are the same as those of Sect. 5.3. However, we can see that activity classes tend to be better served by motion metrics, whereas context classes are based more on Bluetooth, microphone, or network metrics. Finally, the two “Average” lines indicate a common combination in all classes that minimizes the average error rate. For example, the average context category combination is the one with the lowest average error rate for the classes of the category. These two lines are used in the next section to determine users’ classes.

Table 4. Influence of each metric on the recognition of classes. The red to yellow gradient indicates high to low prediction accuracy. Grey indicates that no data was available for the performed activity.

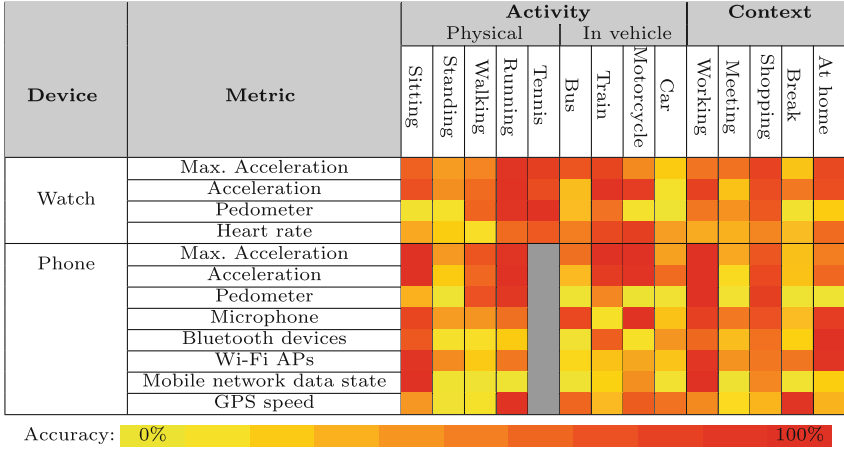


Table 5. Best combinations of metrics for each activity. The “Best accuracy” columns denote the best possible percentage of the test data set which is correctly identified in the training data set. “Average” rows show the best combination for the entire class.

Category	Class	Best combination										Best accuracy				
		Watch		Phone								Phone & Watch	Phone only	Watch only		
		Max. Acceleration	Acceleration	Pedometer	Heart rate	Max. Acceleration	Acceleration	Pedometer	Bluetooth devices	Microphone	Wi-Fi APs				Mobile network data state	GPS speed
Activity	Sitting													99.5	99.0	85.6
	Standing													95.3	81.4	68.8
	Walking													95.8	87.0	73.2
	Running													100	97.4	100
	Tennis													100	-	100
	Bus													88.9	69.5	73.1
	Train													100	96.6	83.9
	Motorcycle													100	97.8	100
	Car													93.2	90.4	77.0
Average													93.2	88.1	81.5	
Context	Working													95.3	95.3	72.3
	Meeting													94.7	85.8	73.4
	Shopping													86.6	76.7	80.6
	Break													95.5	84.9	69.4
	At home													100	99.5	92.4
	Average													90.4	86.6	74.7

5.5 Application Example

To illustrate our conclusions on the data set, we have taken as an example the participant shown in Fig. 2. Each activity and context class is identified using the average combinations (Table 5). The recognition method is applied by progressively comparing the individual user’s data with the data in our full data set using SVM. Figure 4(a) and (b) illustrate the activity and user context recognition respectively, when the user’s data is not included in the full data set.

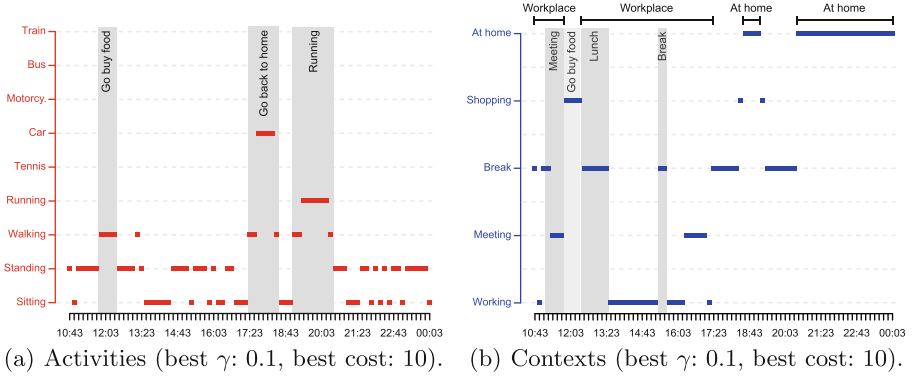


Fig. 4. Detected classes. Grey bars and black text are the main activities and contexts reported by the user.

The participant’s data is divided into ten-minute intervals. For each interval, we calculate the mean and the standard deviation of each metric. The set of values for the participant is small and therefore relatively easy to obtain. Each ten-minute interval consists of 14 values for activity detection and ten values for context recognition. As we can see from the figures and by consulting the participant’s questionnaire, we obtain a very realistic result, which is made possible by the collaboration of all participants and the pooling of their data. In Fig. 4(a), for example, we see that at around 18:00 the participant was driving, and at between about 19:00 and 20:00 he was running. He took a lunch break around noon, which required him to move (walk) to buy food, as confirmed in Fig. 4(b). The same figure also indicates that the participant was in his office most of the afternoon. Some errors are noted around 18:00, where it was detected that the participant was in a shopping centre, where in fact he was in his car. Similar findings are noted around 19:00, when the participant was running. The reason for this is the lack of a corresponding context class, and therefore the closest alternative is indicated.

Figure 5(a) and (b) respectively show the changes in the number of Wi-Fi access points and Bluetooth devices that the participant encounters. It is interesting to compare these figures with the previous ones, because they highlight certain geographic and social characteristics. In Fig. 5(a), for example, there is

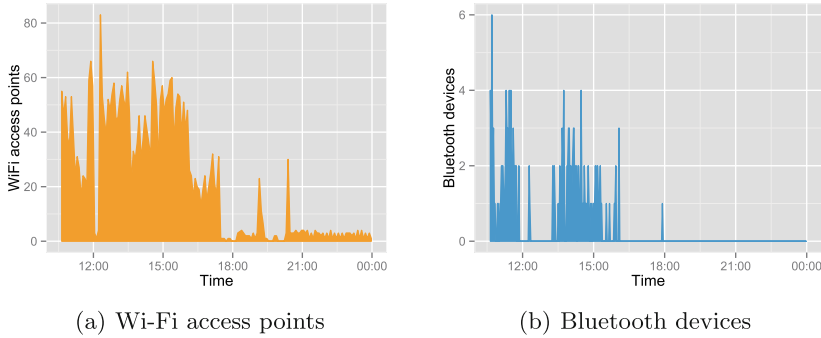


Fig. 5. Geographic and social characteristics.

a huge difference in the number of Wi-Fi access points encountered before and after 18:00, suggesting that the participant visited two major places (in this case, a work environment and a domestic one). It is also interesting to observe the dip around 12:00, which is when the participant visited the shopping centre. The participant’s movement, by car around 18:00 and while running between around 19:00 and 20:00, is also indicated by some slight spikes that we are often associated with travel: the more a participant moves, the more he comes into proximity with different access points. However, these figures do not provide a particularly accurate information base for estimating the participant’s social interactions. To do this, in the following section, we compare activities and social interactions among the participants.

6 Comparing Participants

If the recognition of user activity is an essential step that we can approach with great accuracy, another critical step is to compare several participants. In this section, we introduce novel visual representations, allowing comparison of the 13 participants in the study.

Figure 6 is a 3D plot showing the distribution of types of activity following three different axes. The first reflects the proportion of time the participants were inactive (e.g. sitting). Because the measurements were taken during workdays, the proportion is very high and goes from 63 % (P13) to 90 % (P8). The second reflects the proportion of time the participants were active and were performing an activity (e.g. walking, running). This number distinguishes two categories: those with a sporting activity outside of work and those who are required to move (e.g. to meetings). The third axis reflects the proportion of time the participants were aboard a vehicle. This number is the lowest, and corresponds mainly to journeys between work and home. However, participants such as P10 or P11 have work activities involving frequent trips during the day (e.g. to move from one campus to another). Finally, note that the size of a bubble is proportional to the sum of all acceleration recorded by the watch. Thus, a small bubble indicates

very little sports activity while a larger bubble indicates more frequent, abrupt movement (e.g. running).

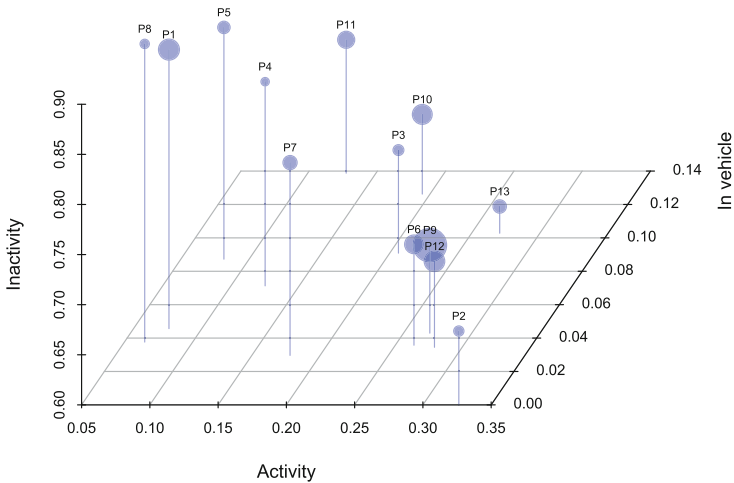


Fig. 6. Comparing activities of the participants.

Figure 7 uses the same principle as the previous figure but is based on three network metrics. First, the average mobile network data state tends to 0 if the mobile phone is connected to a Wi-Fi access point and it tends to 1 if the mobile phone uses cellular data. As the devices are set up only to connect to workplace access points (Table 2), this value is a good indicator of whether the user is more likely to be in the workplace or outside. The number of different access points gives us information about geographic locations visited by the participants. If two people are working in the same place, the participant with the higher value is moving around more and coming into contact with more access points. Finally, the number of distinct devices encountered gives us a measure of the interaction that the participants have. The higher this number, the more devices (a proxy for people) the person has encountered during his or her recording session.

Comparing the two graphs allows us to make some interesting observations. For example, participant P9 seems to perform more physical activity than anyone else, judging from his relatively high activity rate. Moreover, looking at Fig. 7, we find that P9 does not spend much time at the workplace, as he or she encounters the lowest number of access points. Conversely, participant P7 was mainly working during the study and hardly moved at all. Participant P4 is an interesting case, since he or she seems to have been in a vehicle and been in the proximity of a large number of access points. This indicates movement through many public spaces or buildings.

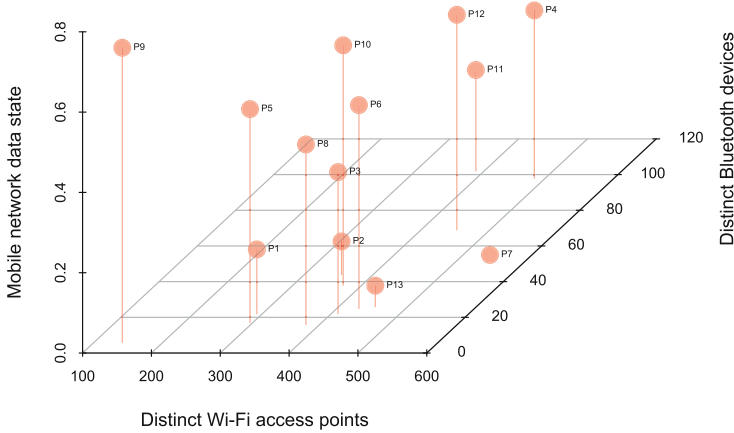


Fig. 7. Comparing interactions of the participants.

7 Conclusion

In this paper, we have described a strategy for recognizing the activities and the contexts within which a user is located. Our results show that using a condensed data set, along with energy-efficient sampling parameters, has the advantage of being easy to use with a classification algorithm such as SVM. Moreover, as such a structure implies lower transmission, harvesting and recording frequencies, it allows energy savings (resulting in an autonomy of about one day using our sensing system). We then showed that using a smartwatch in addition to traditional smartphones leads to better detection accuracy, in particular regarding physical activities such as running (100 % accuracy over our dataset) or walking (95.8 %). In addition, as these wearables are permanently on the user’s wrist, they can detect specific activities without the help of any smartphone (e.g. tennis). Overall, the use of multimodal metrics as advanced feature sets for an SVM model allows the recognition of nine user-defined activities and five contexts, with an average accuracy greater than 90 %. Finally, we presented a new approach that graphically compares the activity and social relations of different users, allowing a better understanding of their behaviour.

The relatively small number of participants and their often vague answers to the questionnaire prevented us from expanding our data set. However, the study suggests great potential for the detection of personal activities if carried out on a wider sample group of users. In future work, in addition to using new devices and extending our energy saving strategy, we plan to carry out similar tests on a larger scale, performing new experiments and/or using public data sets. This will not only allow us to use other learner types and refine our classification model (e.g. adding FFT-based features), but also to accumulate a more extensive activity database that can be used as training set. We also plan to extend our

study to capture user activities and contexts on a weekly basis, which would further help us to recognize patterns and characteristics specific to each user.

References

1. Xu, C., Pathak, P.H., Mohapatra, P.: Finger-writing with smartwatch: a case for finger and hand gesture recognition using smartwatch. In: Proceedings of the 16th International Workshop on Mobile Computing Systems and Applications, pp. 9–14. ACM (2015)
2. Tilenius, S.: Will An App A Day Keep The Doctor Away? The Coming Health Revolution. Forbes CIO Network (2013)
3. Campbell, A.T., Eisenman, S.B., Lane, N.D., Miluzzo, E., Peterson, R.A.: People-centric urban sensing. In: Proceedings of the 2nd annual international workshop on Wireless internet, p. 18. ACM (2006)
4. Burke, J.A., Estrin, D., Hansen, M., Parker, A., Ramanathan, N., Reddy, S., Srivastava, M.B.: Participatory sensing. In: Center for Embedded Network Sensing (2006)
5. Lane, N.D., Miluzzo, E., Lu, H., Peebles, D., Choudhury, T., Campbell, A.T.: A survey of mobile phone sensing. *IEEE Commun. Mag.* **48**(9), 140–150 (2010)
6. Yang, F., Wang, S., Li, S., Pan, G., Huang, R.: MagicWatch: interacting & segueing. In: Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication, pp. 315–318. ACM (2014)
7. Faye, S., Frank, R.: Demo: using wearables to learn from human dynamics. In: Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services, pp. 445–445. ACM (2015)
8. Zheng, X., Ordieres-Meré, J.: Development of a human movement monitoring system based on wearable devices. In: The International Conference on Electronics, Signal Processing and Communication Systems (ESPCO 2014) (2014)
9. Bluetooth, S.: Bluetooth specification version 4.0. In: Bluetooth SIG (2010). <http://www.bluetooth.org/en-us/specification/adopted-specifications>
10. Rodrigues, J.G., Aguiar, A., Barros, J.: SenseMyCity: Crowdsourcing an Urban Sensor. arXiv preprint [arxiv:1412.2070](https://arxiv.org/abs/1412.2070) (2014)
11. Cuervo, E., Balasubramanian, A., Cho, D.-K., Wolman, A., Saroiu, S., Chandra, R., Bahl, P.: MAUI: making smartphones last longer with code offload. In: Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services, pp. 49–62. ACM (2010)
12. Honicky, R., Brewer, E.A., Paulos, E., White, R.: N-smarts: networked suite of mobile atmospheric real-time sensors. In: Proceedings of the Second ACM SIGCOMM Workshop on Networked Systems for Developing Regions, pp. 25–30. ACM (2008)
13. Hussain, S., Bang, J.H., Han, M., Ahmed, M.I., Amin, M.B., Lee, S., Nugent, C., McClean, S., Scotney, B., Parr, G.: Behavior life style analysis for mobile sensory data in cloud computing through MapReduce. *Sensors* **14**(11), 22001–22020 (2014)
14. Han, M., Lee, Y.-K., Lee, S., et al.: Comprehensive context recognizer based on multimodal sensors in a smartphone. *Sensors* **12**(9), 12588–12605 (2012)
15. Lu, H., Pan, W., Lane, N.D., Choudhury, T., Campbell, A.T.: Sound- Sense: scalable sound sensing for people-centric applications on mobile phones. In: Proceedings of the 7th International Conference on Mobile systems, Applications, and Services, pp. 165–178. ACM (2009)

16. Ma, L., Smith, D., Milner, B.: Environmental noise classification for context-aware applications. In: Mařík, V., Štěpánková, O., Retschitzegger, W. (eds.) DEXA 2003. LNCS, vol. 2736, pp. 360–370. Springer, Heidelberg (2003)
17. Sharma, V., Mankodiya, K., De La Torre, F., Zhang, A., Ryan, N., Ton, T.G.N., Gandhi, R., Jain, S.: SPARK: personalized parkinson disease interventions through synergy between a smartphone and a smartwatch. In: Marcus, A. (ed.) DUXU 2014, Part III. LNCS, vol. 8519, pp. 103–114. Springer, Heidelberg (2014)
18. Shin, D., Shin, D., Shin, D.: Ubiquitous health management system with watch-type monitoring device for dementia patients. *J. Appl. Math.* **2014**(2014), Article ID 878741, 8 (2014). <http://dx.doi.org/10.1155/2014/878741>
19. Porzi, L., Messelodi, S., Modena, C.M., Ricci, E.: A smart watch-based gesture recognition system for assisting people with visual impairments. In: Proceedings of the 3rd ACM International Workshop on Interactive Multimedia on Mobile & Portable Devices, pp. 19–24. ACM (2013)
20. He, Z., Liu, Z., Jin, L., Zhen, L.-X., Huang, J.-C.: Weightlessness feature—a novel feature for single tri-axial accelerometer based activity recognition. In: 19th International Conference on Pattern Recognition, ICPR 2008, pp. 1–4. IEEE (2008)
21. Kao, T.-P., Lin, C.-W., Wang, J.-S.: Development of a portable activity detector for daily activity recognition. In: IEEE International Symposium on Industrial Electronics, ISIE 2009, pp. 115–120 (2009)
22. Qian, H., Mao, Y., Xiang, W., Wang, Z.: Recognition of human activities using SVM multi-class classifier. *Pattern Recogn. Lett.* **31**(2), 100–111 (2010)
23. Wu, J., Pan, G., Zhang, D., Qi, G., Li, Shijian: Gesture recognition with 3-D accelerometer. In: Zhang, D., Portmann, M., Tan, A.-H., Indulska, J. (eds.) UIC 2009. LNCS, vol. 5585, pp. 25–38. Springer, Heidelberg (2009)
24. Jain, A.K.: Data clustering: 50 years beyond K-means. *Pattern Recogn. Lett.* **31**(8), 651–666 (2010)
25. Bishop, C.M.: *Pattern Recognition and Machine Learning*. Springer, New York (2006)
26. Dimitriadou, E., Hornik, K., Leisch, F., Meyer, D., Weingessel, A.: e1071: Misc Functions of the Department of Statistics (e1071), TU Wien (2011). <http://CRAN.R-project.org/package=e1071>
27. Bergstra, J., Bengio, Y.: Random search for hyper-parameter optimization. *J. Mach. Learn. Res.* **13**(1), 281–305 (2012)