Toward a Characterization of Human Activities using Smart Devices: A Micro/Macro Approach

Sébastien Faye, Nicolas Louveton, Gabriela Gheorghe, Thomas Engel
Interdisciplinary Centre for Security, Reliability and Trust (SnT)
University of Luxembourg
4 rue Alphonse Weicker, L-2721 Luxembourg, Luxembourg
Email: {firstname.lastname}@uni.lu

***Abstract*—The emergence of new connected devices has opened up new opportunities and allowed to imagine concepts that bring computer sciences and social sciences closer together. In particular, today's increasingly sophisticated miniature sensors allow to track and understand human activities and behavior with a great precision. Taking different approaches and perspectives, we use in this paper smartwatches and smartglasses to explore these behaviors and show that these objects, considered by many as gadgets, have an important role to play in understanding the lives of individuals. The main objective of this work is to introduce two new scales of activity detection, which lacks a formal and consistent definition in the literature. First, we propose a model that precisely detects and interprets movements made by a person wearing smart devices. Then, we use this model to show different interactions between those *micro-activities* and bigger chunks of behaviors we call *macro-activities*. Using a new concept based on 3D visualization, we finally show that combining those two scales and using a limited dataset, it is possible to distinguish between different individuals when they are performing very similar activities. The findings of this study lead the way to enhanced user profiling.**

***Index Terms*—Activity and Mobility Detection, Sensing Systems, Mobile Computing, User Profiling.**

I. Introduction

In recent years, the growing availability and falling cost of connected objects embedded with software and sensors has opened up a world of opportunities for new applications. Apart from smartphones, which are not all the time on the user, these objects include a wide range of ultra-portable (or *wearable*) devices that constantly interact with the body movements and with the IT environment (e.g. using Bluetooth of Wi-Fi). These wearables have much in common with ubiquitous computing both in the field of research and in terms of functionality [1]. They can precisely detect a range of events and can be precise, responsive and permanently operational [2].

Recently, the arrival on the market of major players like Apple, Google and Microsoft popularized smartwatches and smartglasses and facilitated the development and widespread adoption of sensing applications (e.g. with Android Wear), opening doors in many areas including sport and personal monitoring. Moreover, data from activity sensors and sports watches, usually proprietary to device manufacturers [3], can be accessed via services such as Apple Health and Google Fit. Technological advances have led to the integration of sensors that can produce results equal to those of specialized experimental devices. [4] shows for example that movements of the arms, the hands and possibly the fingers, generate energy strong enough to be picked up by the accelerometer and the gyroscope of a smartwatch with 98% precision.

Detecting and identifying user activity has undergone such extensive research that it is already being integrated into many commercial products. From the fitness tracking market, Google Fit App claims for example to detect basic activities such as walking, running or biking. However, the concept of user activity lacks of a formal and consistent definition across the state of the art or the commercial applications. Commercial approaches focus on activity understood as physiological and motor activity. While physiological and motor activity can easily be captured by wearable devices through commercial sensors, they point to a limited area of the human activity spectrum. In that respect, Newell [5] described a range of behavioral levels (called bands) distinguished by the frequency of occurrence of a specific behavior: biological, cognitive, rational and social bands (see Table I for details).

	Newell scales	Level	Duration (s)	New scales	Example
Individual	Biological	Cell / cell groups	$10^{-4:-2}$	Micro	Neuron spike
	Cognitive	Physiological response or motor act	$10^{-1:1}$		Jumping, taking a step, heart beat
	Rational	Hierarchy of actions towards a goal	$10^{2:4}$	Macro	Walking or running toward a direction
Group	Social	Coordinated goal	$10^{5:7}$		Sport competition

TABLE I: Activity scales

In the state of the art, there is no distinction between those scales. Moreover, activity detection frequently refers to the cognitive band only (activities lasting from split-second to a few seconds). Those activities are elementary actions that support the achievement of larger scale goals. For this reason, we propose two new concepts in this work in order to support a better understanding of user activities: on the one side, **micro-activities** are activities detectable from wearable sensors measurements that relate to activities located in the cognitive or even in the upper-bound biological bands; on the other side, **macro-activities** are activities located in the rational or in the lower-bound social bands. In our view, users reason in macro-activities (e.g. physical exercise). Macro-activities are meaningful to users because they can be directly correlated with goals. Such goals can be decomposed into activities that individually are less meaningful: micro-activities (e.g. walk steps, squats) do not carry value by themselves.

This paper reports on experiments where we used two smart devices for extracting micro- and macro-activity information. After introducing related work in Section II and our methodology in Section III, we show in Section IV how smartglass devices could be used to determine head gestures along with a long-run visualization of those gestures. A long-run visualization of directional head movement in combination with social ones might be a possible way to understand group interactions in a given environment. In contrast, for the reporting of Section V we used smartwatches to classify users in different locomotion activities using a three sensor approach and a 3D visualization of the data. To summarize, this work has the following objectives: (1) to present two concepts, i.e. micro- and macro-activities; (2) to characterize and compare those micro- and macro-activities using two smart devices to collect data; (3) to present future ways of using these new devices to improve our understanding of human activities.

II. Related Work

A. Sensing Systems

The use of smart devices as key elements in a sensing platform has been discussed for many years, in both industrial and research communities [6]. Wearables such as smartwatches and smartglasses have their place in this ecosystem and can open up new perspectives. By combining these devices, a large amount of data can be obtained from sensors such as GPS, accelerometer, gyroscope, magnetometer and most recently even barometer, temperature or heart rate monitor, as well as interactions with Bluetooth, WiFi, NFC or cellular [7].

In most cases, smart devices are connected via Bluetooth Low Energy (BLE) [8], a relatively new technology that has been standardized under the Bluetooth 4.0 specification [9]. Detection can be as opportunistic [10] as it is participatory [11]. In the case of a participatory system, incentives are an important aspect as the user is involved in the data collection process: they can include services, games or personal benefits [7]. In the case of an opportunistic system, a lot of applications are based on crowd-sensing to aggregate data from sensors [12].

Finally, note that these kind of ecosystems can address three levels of detection [6]. At individual level, where detection and data processing are targeted at and for the individual (or perhaps some authorized person), as in the case of certain applications connected to sport [13]. At group level, where individual participants with devices share a goal or a common interest, when there is an element of trust within the group [14]. At community level, with many participants [15]. At this level, when people are strangers and don't have the same trust in one another, it is important to define rigorous systems for the protection of personal privacy.

B. Activity and Context Detection

As introduced, the direct activity and movements of the user can be reliably deduced by smart devices [4], [16]. FDSVM [17] (Frame-based Descriptor and multi-class SVM) is an approach that can classify a large variety of gestures using a three-axis accelerometer. It is based on SVM [18] (Support Vector Machine), a set of statistical learning methods intended to resolve problems of discrimination and regression. [19] present a pointing device that detects movement with FDSVM and also takes social interaction into account. Other algorithms are also known and are based on the DTW approach [20] (Dynamic Time Warping) or even HMM [21] (Hidden Markov Model).

Detecting motion and different types of activity (walking, running, traveling in a car, etc.) is also possible using a smartphone alone [22]. However, this choice is not the most relevant, given that the user doesn't always have the smartphone with him. Moreover, studies in the University of Washington – Intel Mobile Sensing Platform [23] show a real interest in activity detection using new sensors, still little used on the market (e.g. walking on a wet floor using humidity sensors). [24] studies patients with mental disorders and uses smartwatches to help quantify the exercise and the amount of sunlight wearers received. The measured data is sent to a server where it can be analyzed and consulted by doctors.

Finally, smart devices can detect and understand the context (or environment) in which a user finds himself. [25] uses data from the accelerometer, sound, GPS and WiFi signals to classify its activities. It also monitors ambient noise. By continuously recording sound, it is possible to identify the contexts of a person's life, whether having a conversation, sitting in the office, walking out on the street, or even making coffee [26]. EmotionSense [27] is a platform that uses data collected by smartphones for social psychology studies, detecting activities but also verbal interaction and proximity to others. [28] automatically classifies personal events, and automatically shares the results of its analysis on social networks.

III. Methodology

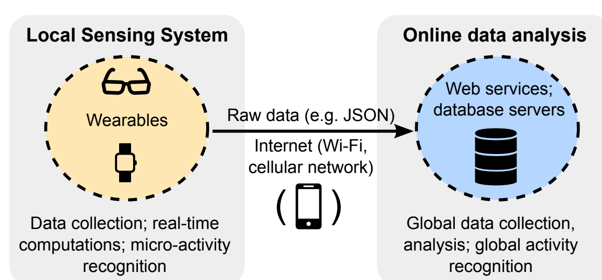

Fig. 1: Overview of our architecture

The architecture used as part of our studies is illustrated in Figure 1. It is composed of two main parts. First, the sensing system is based on one or more smart devices with the objective of collecting metrics when users wear them. This is done in particular through an Android application that we have developed independently for smartglasses and smartwatches, which usually need a smartphone to access the Internet. Then, an online platform is responsible for data storage and overall processing. The devices used as part of our experiments are detailed in Table II.

Devices	RAM / Storage	CPU	Network Interfaces	Main sensors
EPSON Moverio BT-200	1 GB / 8 GB	Dual-core 1.2 GHz	GPS, 802.11 b/g/n, Bluetooth 3.0	Accelerometer, gyroscope, compass, microphone.
Samsung Gear Live	512 MB / 4 GB	Quad-core 1.2 GHz	Bluetooth 4.0	Heart rate, pedometer, accelerometer, gyroscope, compass.

TABLE II: Specification of the devices

In this study, we use smartglasses for detecting micro-activities, and smartwatches for detecting macro-activities.

Indeed, as micro-activities are short actions, they do not need a long period of data collection. Smartglasses are generally less comfortable than smartwatches, which can be worn over a longer time. In addition, micro-activities are more sensitive to movements, i.e. we need to collect data at a high rate. Smartglasses are for studying micro-activities because head movements are generally less important than arm movements with respect to goal achievement. The sensitivity of the system is therefore important. We used the Epson MOVERIO BT-200 smartglasses for our experiments because of their advanced and open operating system (Android). However, the methods described in this paper are applicable to any type of device running on Android 4+.

Smartwatches were then selected for the detection of macro-activities, as they are comfortable and integrate additional sensors to detect specific activities (e.g. being in public transportation detecting nearby devices and access points). Moreover, these devices have a good battery capacity and can detect the IT environment using Bluetooth 4.0.

IV. Detecting Micro-Activities

In this section, we use Dynamic Time Warping (DTW) to detect micro-activities. DTW is one of the most popular algorithms that measures the similarities between two time sequences of different speed and duration [29]. Using this method on smartglasses and based on a small dataset collected over three participants (Sec. IV-C), we then study the advantages and disadvantages of using those micro-activities to understand the human behavior and mobility.

A. Dynamic Time Warping

DTW aligns two time series in order to minimize the cumulative distance d between each of the data points.

Suppose we have two one-dimensional time series: $X = (x_1, ..., x_{|X|})$ and $Y = (y_1, ..., y_{|Y|})$. DTW provides a cost, which is calculated through a two-dimensional cost matrix C ($|X|$ by $|Y|$). Each cell represents the minimum cost accumulated between the time series X and Y to the position of that cell, and is calculated by:

$$C_{(i,j)} = D(x_i, y_i) + min\{C_{(i,j-1)}, C_{(i-1,j)}, C_{(i-1,j-1)}\} \quad (1)$$

which represents the distance between point i of series X and point j of series Y, of the minimum accumulated distance from the previous three cells that surround cell i, j (the cell at the top, left and diagonally). D is a distance function, generally Euclidean:

$$D(x_i, y_i) = (x_i - y_i)^2 \quad (2)$$

When the matrix is full, the minimum normalized distance between X and Y is obtained by taking the value of the last cell:

$$DTW(X, Y) = D(|x|, |y|) \quad (3)$$

Figure 2 shows an example of the alignment of multiple time series, taking the three axes of the gyroscope and describing a “Yes”, nodded while wearing smartglasses.

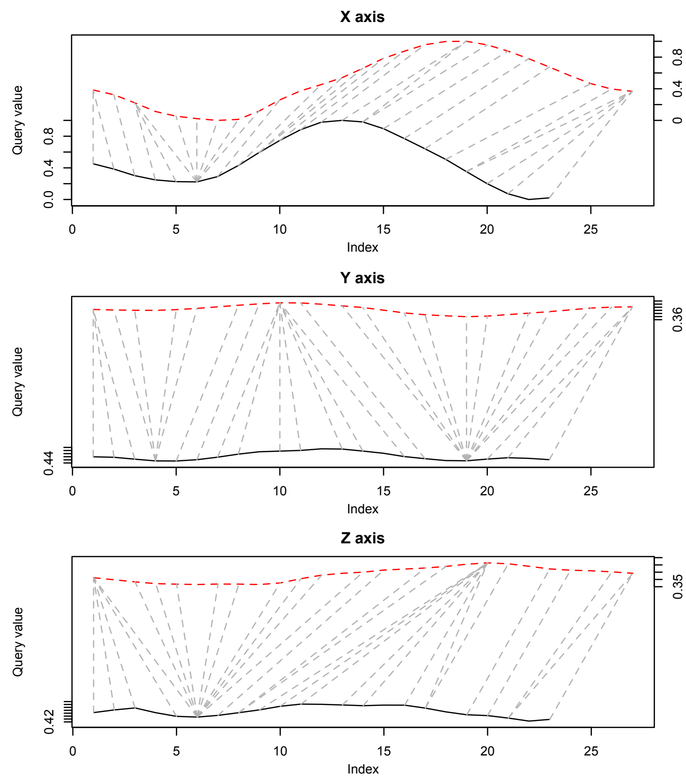

Fig. 2: Alignment of multiple time series describing a “Yes”

B. Considering multiple dimensions

In order to perform motion detection and recognition, we set up a program to record a training dataset composed of multiple time series identified by a label (i.e. the name of the movement being performed – e.g. nodding head).

Using the three axes of the gyroscope as the basis for detecting and registering head movements, we then defined two groups and a total of seven movements, i.e. **social**: yes, no, don’t know (which corresponds to a shrug and a slight tilt of the head) and **direction**: up, bottom, right, left.

After each movement led one or more times, the detection of a movement is done simply by comparison with the training dataset. In order to treat the three dimensions represented by the three axes of the gyroscope, we use a Euclidean normalization (L2 norm) to reduce each movement to a one-dimensional series $\sqrt{\alpha^2 + \beta^2 + \gamma^2}$, with α, β and γ axes of the gyroscope.

C. Experiments

We have implemented our method on smartglasses whose choice was justified in the previous section. The setting is of an Android application that allows (1) to record a training dataset consisting of several movements defined by the user, and (2) to return the movement and the distance determined using DTW that match a current movement.

To record the user's movements, the application automatically detects sudden changes of the gyroscope, and stores or classifies the corresponding time series. Two other methods were also implemented, one based on a time slot (e.g. registration or classification every 5 seconds) and the other based on a button that the user must push to start or stop recording.

Our experiments were based on three different participants. The first recorded the movements described in Sec. IV-B for a training dataset. The second wore glasses for 2.8 h at different places (work, public transport, city center, etc.), which caused the detection of 1,863 movements (i.e. one every 5.4 seconds). The third wore them for 3.6 h and was all the time in the same room, which caused the detection of 1,689 movements (i.e. one every 7.65 seconds). It seems logical that an inactive person does less movements that a person who is moving.

D. Results

Figure 3 shows the number of movements (described in Sec. IV-B) detected in both groups for the second participant, whose behavior seems more interesting than the third. The red dots show the difference between group A (direction) and group B (social). After consulting a diary activity[1] provided by the participant, we can draw some conclusions. Around 2,000 and 8,000 seconds, the participant was sitting in a bar, which caused more social movements. Between 3,000 and 5,000 seconds, the participant was in the street walking with a variable speed. The peak of attention seems longer between 3,000 and 4,000 seconds, when the participant was in the middle of a fun-fair. Finally, before 2,000 and around 6,500 seconds, the participant was not wearing the hardware, which explains the lack of movements. This allows to understand more precisely the type of interaction performed by a user, without having to use expensive methods (e.g. image processing).

Figure 4 shows the distribution of distances computed for each registered gesture. The graph shows three categories of gestures that are detected each with different accuracies. This means that the training dataset is not representative of all types of movements that can be performed, and even less of different people that can make these movements. Indeed, each individual has his own way of behaving, and this scale study is too precise to easily identify real differences.

[1] Available online: http://swipe.sfaye.com/mobiworld16/diary-activity.pdf

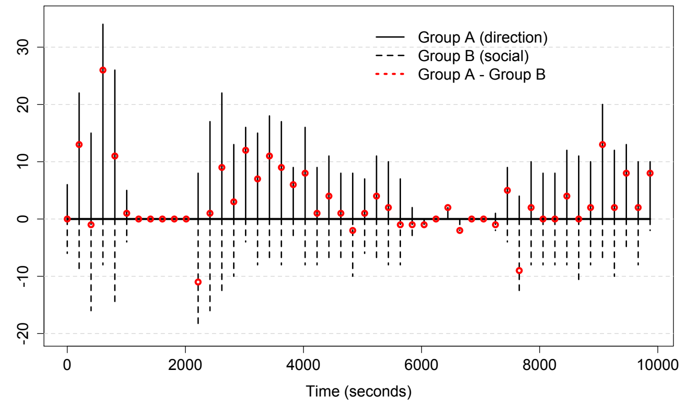

Fig. 3: Micro-activity detection timeline

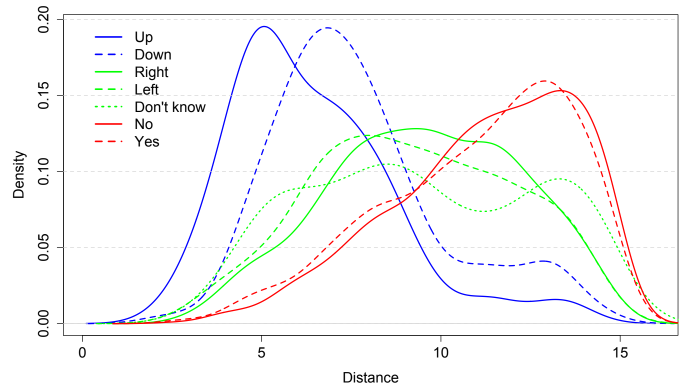

Fig. 4: Micro-activity detection distribution. The distance is computed according to the formula 2 used for one-dimensional series (Section IV-B).

As just presented, the combined detection of several micro-activities can describe moments in the user's life and forms macro-activities. This approach can be useful to understand how small physical actions can compose into complex and goal-directed actions. Furthermore, this approach has the potential to consider other metrics, such as the ambient sound or the orientation of the smartglasses to detect additional contextual information. In the next section, we use micro-activities in order to distinguish between users performing the same given macro-activity.

V. Using Macro-Activities to Classify Users

In this section, as discussed above, we assume that we have enough tools and technologies to easily detect macro-activities. In [30], for example, we proved that it is possible to use a combination of different sensors from a smartphone and a smartwatch to describe physical and social behaviors of different users. Other algorithms are also relevant, such as those integrated to Android to detect physical moves of a user.

In order to go further in trying to compare and classify users doing similar activities, our idea is to combine different

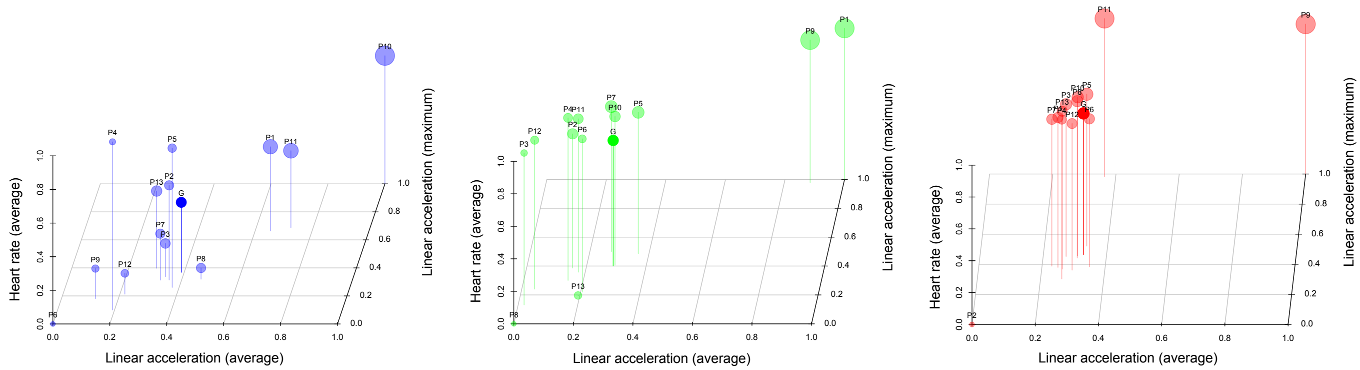

(a) Inactivity (b) Physical activity (c) Transportation

Fig. 5: Three aspects of the user's everyday life

macro-activities with micro-activity metrics, i.e. raw data from sensors. To this extent, we introduce novel 3D visual representations, allowing the comparison of different users with similar activities. Moreover, we rely on a dataset obtained from 13 participants we have collected using a smartwatch, a smartphone and the SWIPE open-source application [31] (i.e. 157.2 h of recording). More details can be found in [30].

First of all, we choose three metrics to represent different aspects of a user's everyday life. We avoid using metrics such as GPS or environmental data because we want to focus on the behavior of the user, i.e. his movements as well as physiological data. For this reason, we choose to use (1) the **average heart rate** of the user from the smartwatch (recorded every 60 seconds), (2) the **average linear acceleration** computed on the smartwatch (recorded every 30 seconds) and (3) the **maximum linear acceleration** on the smartwatch to detect sudden movements (recorded every 30 seconds). The linear acceleration can be described as the rate of change of velocity of an object. It is computed over the three axes of the accelerometer: $\sqrt{x^2 + y^2 + z^2}$.

These metrics are then computed for three complementary aspects as shown in Figure 5, namely **inactivity**, **physical activity** and **transportation**. Figure 5(a) is a 3D representation of the three normalized metrics computed when the users was in the "sitting" position. Figure 5(b) represents the users when they were in a "walking" activity. In both cases, we can see a relation between the maximum and the average linear acceleration, and a good distribution of all users. This means that each user has his own way of moving, and it is visible easily using this kind of graph. Finally, figure 5(c) shows when the users was in a vehicle. The distribution is different from the two others as we clearly see one big cluster, showing participants that use their own car. The two other participants were using public transportation (bus, train).

Figure 6 gives an idea of how we can compute a profile for each user. The idea is to combine the three previous graphs and to compute three distinct indexes: activity index, inactivity index and transportation index. Each index is normalized between 0 and 1 and is computed depending on the gravity center of the graphs, which is a common value between all data points: it is used to compute the distance of each user compared to the average behavior. For example, the activity index of P11 is the absolute value of the distance between the gravity center of all other participants and the P11 point in Figure 5(b). We can see for example than the profile of P1 seems to be close to the one of P8. This example is just an introduction to the possibilities we have using such aggregated graphical methods and can lead to user profiling and enhanced personalized services as we discuss in the next section.

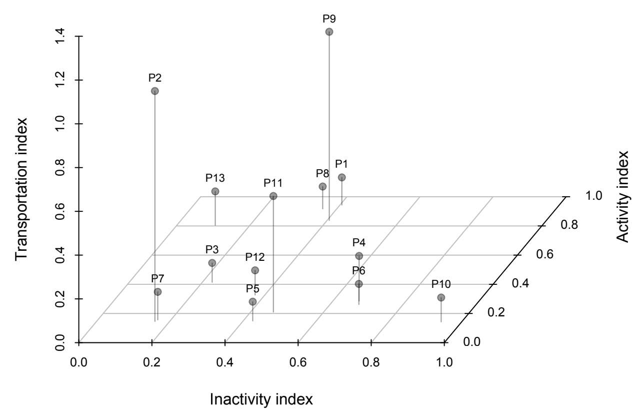

Fig. 6: Global index

VI. Conclusion and Perspectives

In this paper, we presented two activity detection scales and studied their interactions in order to improve the understanding of human behavior and mobility. Using new visualization methods drawn from smartglasses and smartwatches data, we show that micro- and macro-activities could be analyzed together for a better understanding of the user behavior: on the one hand, combining micro-activities can provide a better understanding of user activity overall (like in Figure 3); on the other hand, within a given macro-activity, individuals could be identified using micro-activity patterns (see Figure 5).

In our point of view, the originality of combining micro- and macro-activity levels of behaviors will lead to two types of innovative applications and services. First, if a gesture is the unit of user's daily routine as a word is the unit of a sentence, combining different levels of activity will allow for a better understanding of the user's goals which in turn will improve tailoring applications and services to his/her needs. Then, having a complete scope of a user's daily routine will make it easier to match user content or service recommendations by using preferences of users who share the same routine. Such an approach can rise challenges in terms of the trade-off between privacy constraints and profiling accuracy.

Beyond detection and classification, sensing activities and environments may open possibilities of long-term life-logging and memory augmentation. More specifically, long-term memory mostly works based on contextual cues and priming that actually activate more abstract, verbal memories about facts and self. A good illustration of this is Proust's "episode of the madeleine". What if a wearable device could sense, store and (re-)activate context automatically and on a large scale? Applications are ranging from personal information management, social sharing of stories and activities, and even medical applications related to supporting memory or detecting disease early-on. Finally, it is important to note that those applications rise important issues in terms of privacy: how to improve user awareness of what is currently being logged, how it is processed or shared, and how the user may tune the sensitiveness of context sensing.

Acknowledgment

This work has been performed as part of the eGLASSES project, which is partially funded by NCBiR, FWF, SNSF, ANR and FNR under the framework of the ERA-NET CHIST-ERA II.

References

[1] B. J. Rhodes, N. Minar, and J. Weaver, "Wearable computing meets ubiquitous computing: Reaping the best of both worlds," in *Wearable Computers, 1999. Digest of Papers. The Third International Symposium on*. IEEE, 1999, pp. 141–149.

[2] G. Bieber, M. Haescher, and M. Vahl, "Sensor requirements for activity recognition on smart watches," in *Proceedings of the 6th International Conference on PErvasive Technologies Related to Assistive Environments*. ACM, 2013, p. 67.

[3] J. Westlin and T. H. Laine, "Manysense: An extensible and accessible middleware for consumer-oriented heterogeneous body sensor networks," *International Journal of Distributed Sensor Networks*, vol. 2014, 2014.

[4] P. H. P. Chao Xu and P. Mohapatra, "Finger-writing with smartwatch: A case for finger and hand gesture recognition using smartwatch," in *HotMobile 2015. ACM*. ACM, 2015.

[5] A. Newell, "Unified theories of cognitioncambridge university press," *Cambridge, MA*, 1990.

[6] N. D. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, and A. T. Campbell, "A survey of mobile phone sensing," *Communications Magazine, IEEE*, vol. 48, no. 9, pp. 140–150, 2010.

[7] J. G. Rodrigues, A. Aguiar, and J. Barros, "SenseMyCity: Crowdsourcing an Urban Sensor," *arXiv preprint arXiv:1412.2070*, 2014.

[8] K.-H. Kim, M.-Y. Jeon, J.-Y. Lee, J.-H. Jeong, and G.-M. Jeong, "A study on the app development using sensor signals from smartphone and smart watch," *Advanced Science and Technology Letters*, vol. 62, 2014.

[9] S. Bluetooth, "Bluetooth specification version 4.0," *Bluetooth SIG*, 2010.

[10] A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, and R. A. Peterson, "People-centric urban sensing," in *Proceedings of the 2nd annual international workshop on Wireless internet*. ACM, 2006, p. 18.

[11] J. A. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. B. Srivastava, "Participatory sensing," *Center for Embedded Network Sensing*, 2006.

[12] G. Chatzimilioudis, A. Konstantinidis, C. Laoudias, and D. Zeinalipour-Yazti, "Crowdsourcing with smartphones," *Internet Computing, IEEE*, vol. 16, no. 5, pp. 36–44, 2012.

[13] S. Consolvo, D. W. McDonald, T. Toscos, M. Y. Chen, J. Froehlich, B. Harrison, P. Klasnja, A. LaMarca, L. LeGrand, R. Libby *et al.*, "Activity sensing in the wild: a field trial of ubifit garden," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2008, pp. 1797–1806.

[14] CENS/UCLA, "Participatory sensing / urban sensing projects." [Online]. Available: http://research.cens.ucla.edu/

[15] M. Millenium, "Mobile millenium: using cell phones as mobile traffic sensors," *UC Berkeley College of Enginerring, CCIT, Caltrans, DOT, Nokia, NAVTEQ*, 2008.

[16] O. D. Lara and M. A. Labrador, "A survey on human activity recognition using wearable sensors," *Communications Surveys & Tutorials, IEEE*, vol. 15, no. 3, pp. 1192–1209, 2013.

[17] J. Wu, G. Pan, D. Zhang, G. Qi, and S. Li, "Gesture recognition with a 3-d accelerometer," in *Ubiquitous intelligence and computing*. Springer, 2009, pp. 25–38.

[18] N. Cristianini and J. Shawe-Taylor, *An introduction to support vector machines and other kernel-based learning methods*. Cambridge university press, 2000.

[19] F. Yang, S. Wang, S. Li, G. Pan, and R. Huang, "Magicwatch: interacting & segueing," in *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*. ACM, 2014, pp. 315–318.

[20] G. Niezen and G. P. Hancke, "Gesture recognition as ubiquitous input for mobile phones," in *International Workshop on Devices that Alter Perception (DAP08), conjunction with Ubicomp08*. Citeseer, 2008.

[21] T. Schlömer, B. Poppinga, N. Henze, and S. Boll, "Gesture recognition with a wii controller," in *Proceedings of the 2nd international conference on Tangible and embedded interaction*. ACM, 2008, pp. 11–14.

[22] J. Yang, "Toward physical activity diary: motion recognition using simple acceleration features with mobile phones," in *Proceedings of the 1st international workshop on Interactive multimedia for consumer electronics*. ACM, 2009, pp. 1–10.

[23] T. Choudhury, S. Consolvo, B. Harrison, J. Hightower, A. LaMarca, L. LeGrand, A. Rahimi, A. Rea, G. Bordello, B. Hemingway *et al.*, "The mobile sensing platform: An embedded activity recognition system," *Pervasive Computing, IEEE*, vol. 7, no. 2, pp. 32–41, 2008.

[24] D. Shin, D. Shin, and D. Shin, "Ubiquitous health management system with watch-type monitoring device for dementia patients," *Journal of Applied Mathematics*, vol. 2014, 2014.

[25] M. Han, Y.-K. Lee, S. Lee *et al.*, "Comprehensive context recognizer based on multimodal sensors in a smartphone," *Sensors*, vol. 12, no. 9, pp. 12 588–12 605, 2012.

[26] H. Lu, W. Pan, N. D. Lane, T. Choudhury, and A. T. Campbell, "Soundsense: scalable sound sensing for people-centric applications on mobile phones," in *Proceedings of the 7th international conference on Mobile systems, applications, and services*. ACM, 2009, pp. 165–178.

[27] K. K. Rachuri, M. Musolesi, C. Mascolo, P. J. Rentfrow, C. Longworth, and A. Aucinas, "Emotionsense: a mobile phones based adaptive platform for experimental social psychology research," in *Proceedings of the 12th ACM international conference on Ubiquitous computing*. ACM, 2010, pp. 281–290.

[28] E. Miluzzo, N. D. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. B. Eisenman, X. Zheng, and A. T. Campbell, "Sensing meets mobile social networks: the design, implementation and evaluation of the cenceme application," in *Proceedings of the 6th ACM conference on Embedded network sensor systems*. ACM, 2008, pp. 337–350.

[29] E. Keogh and C. A. Ratanamahatana, "Exact indexing of dynamic time warping," *Knowledge and information systems*, vol. 7, no. 3, pp. 358–386, 2005.

[30] S. Faye, R. Frank, and T. Engel, "Adaptive activity and context recognition using multimodal sensors in smart devices," in *The 7th International Conference on Mobile Computing, Applications and Services (MobiCASE'15)*, Berlin, Germany, Oct. 2015.

[31] ——, "Swipe: Monitoring human dynamics using smart devices (demo)," in *The 7th International Conference on Mobile Computing, Applications and Services (MobiCASE'15)*, Berlin, Germany, Oct. 2015.