Behavior modeling and recognition methods to facilitate transitions between application-specific personalized assistance systems

Arun Ramakrishnan[1], Zubair Bhatti[1], Davy Preuveneers[1], Yolande Berbers[1], Aliaksei Andrushevich[2], Rolf Kistler[2], and Alexander Klapproth[2]

[1] IBBT-DistriNet, KU Leuven, Belgium
{firstname.lastname}@cs.kuleuven.be
[2] iHomeLab, Lucerne University of Applied Sciences and Arts, Switzerland
{firstname.lastname}@hslu.ch

Abstract. Activity recognition mandates complex sensor fusion processing. Many contributions in the literature focus on improving the recognition accuracy of a limited set of activities or the efficiency of the algorithms. However, there is little work on how to dynamically adapt the activity recognition techniques when human activity goes beyond the observational and functional borders of one application-specific personalized assistance system. We present tool support to model transitions between activities, and a modular distributed framework of human activity recognition components with support for analyzing resource and recognition trade-offs for different deployments and configurations.

Keywords: Behavior and activity recognition, smart home and health

1 Introduction

The recognition of human activities based on various environment/user aware sensors and associated context information underpins the user-centric paradigm of Ambient Intelligence (AmI). Recent advances in the field of micro-electronics have unfolded the possibility of integrating multiple miniature, cost effective sensors along with standard communication and processing units into single sensing modules that can be effortlessly carried by the user or easily integrated into existing and new construction facilities.

A prerequisite to developing truly supportive and personalized systems is being able to recognize and anticipate typical human behavior and intent in a variety of different contexts [7, 3]. However, the unpredictability of human behavior, the unanticipated circumstances of execution and a growing heterogeneity of future operational environments impose significant development challenges.

In this work, we focus on personalized assistance systems, with examples of domain-specific informational assistance system families including access systems of office buildings, automatic accounting of working time, managing medicine prescriptions and distributions in hospitals, etc. The biggest challenge is however in transitioning and information exchange between those assistance systems when different activities need to be recognized. These recognition techniques might rely on the same sensors or algorithms (on different sensors), but both may need to be fine-tuned at runtime to the new context. Therefore, we need mechanisms to transfer this meta-information from one application to another one covering different activity subsets. Our answer to this challenge is a framework that relies on (1) hybrid behavior modeling that links typical activities in a given situation with different recognition techniques and on (2) an adaptable and modular distributed framework that optimizes the processing and communication of the sensor data in a way that respects the resource constraints of the software and hardware components involved.

2 Related work

A detailed review on the state-of-the-art on activity recognition [2, 3, 8, 4] is beyond the scope of this paper. As our experiments focus on resource trade-offs, we mainly highlight works that looked at device and system energy efficiency as a key concern. A first approach [15] looks at selecting the proper sensors. Accurate user activity prediction needs continuous sampling and the authors propose a method to select an optimal set of sensors at run-time. A similar approach was suggested in [9], arguing that certain sensors are more power consuming than others, with the authors favoring dynamically switching on certain high-cost sensors. Another approach is to adapt the sampling frequency. [13] demonstrates a non-linear relation between sampling frequency and energy consumption especially when frequency domain features (e.g. entropy) are being calculated. It also suggests activity dependent optimal sampling frequencies for mobile based accelerometer sensing, and adapting the window size as the authors found a linear relation between window size and energy consumption on mobile based activity recognition. One can also adapt the features being extracted and selected from the sensor data stream. [1] extensively studied the influence of selected features on classification accuracy and recall for wearable sensors (using 5 accelerometers) and concluded that sometimes fewer features can be more efficient without compensating classification accuracy. Communication efficiency can also significantly impact power consumption, as accelerometer based activity recognition requires high precision data and hence high sample frequencies (up to 100Hz). Techniques to optimize the communication and processing of large amounts of data for wearable and other wireless sensors was the research objective in [12, 11, 6]. Our work focuses on analyzing the trade-offs between all these concerns.

3 Smart home and health motivating use cases

In our activity recognition experiments we compare two *[Anon EU proj]* use cases taken from the smart home and smart health domain. A first activity we aim to recognize is taken from Ambient Assisted Living scenarios, i.e. *detecting a fall* [10, 14, 5]. The first technique only uses a triaxial accelerometer running at 50Hz and looking for patterns of interest through continuous feature extraction

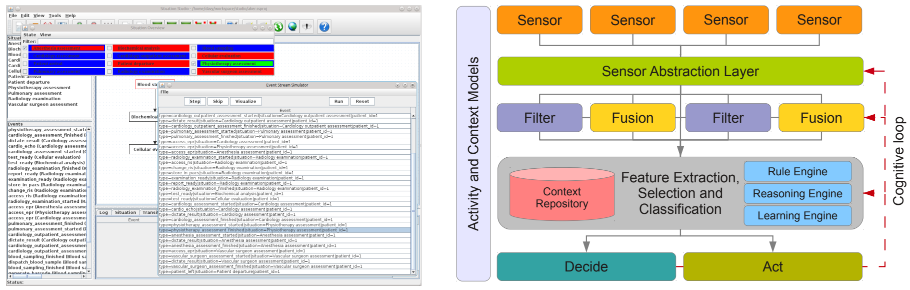

(a) Situation Studio

(b) Human activity recognition architecture

Fig. 1: Activity transition modeling tool (a) and architectural overview (b)

and selection using a high pass FIR filter to take out the gravity component as a fall is characterized by a dynamic acceleration in a small time window (usually less than a second). Our second approach combines the accelerometer with a barometric pressure sensor *[Anon ref]*. A sudden acceleration triggers a sampling of the pressure sensor to detect the current altitude. After a predefined or automatically calculated time delay, another sample is taken. If a significant difference in height is detected, we assume the person has fallen and was not able to get up. This technique does not require complex feature extraction from the accelerometer and pressure sensors.

We use the same accelerometer in a personalized diabetes assistant *[Anon ref]* to *track the exercise level of physical activities* and offer decision support on medicine intake based on past occurrences. We analyze the sensor data stream for activities like walking and running to estimate the calories consumed. This requires feature extraction and selection in both the time and the frequency domain. Other differences with the previous use case are that data should be stored for future reference and comparison, the sampling rate is lower, and the sliding window for signal analysis is several seconds long for better accuracy.

4 A multiple applications approach in behavior modeling and activity recognition

Each of the feature extraction and analysis building blocks are developed as separate components that can be deployed on a sensor mote or mobile running our middleware *[Anon ref]*. In general, the overall distributed architecture is depicted in Figure 1(b). The set of sensors, filters, aggregators, classification and learning components can be deployed, composed and configured dynamically at runtime, depending on the activities to be recognized and corresponding resource trade-offs for wireless communication, computation and memory.

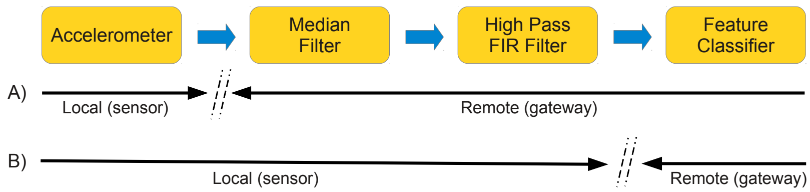

Fig. 2: Two deployment scenarios for feature extraction and selection components

4.1 Modeling and use of contextual domain knowledge

Many works focus on a limited set of activities and validate the accuracy of their approach with the implicit assumption that the activity of interest is taking place. One hardly finds numbers about false positives or false negatives. In our approach, we instead use the cognitive loop in our architecture (see Figure 1(b)) to infer the most probable activities given the current time and location and initiate the corresponding activity recognition techniques. We explicitly consider situations where techniques could lead to false positives, etc. For example, the fall detection with the barometric pressure might detect a false positive when going down the stairs, because with each step the accelerometer triggers the pressure sensor and the latter detects a lower altitude. However, one can also fall down the stairs. All of these interrelationships between different kinds of contexts and activities and corresponding recognition techniques are modeled with our Situation Studio *[Anon ref]*. This tool (see Figure 1(a)) borrows concepts from workflow modeling languages, and represents situations that evolve from one to the next through constrained sequential and parallel transitions. For each of them, we identity the contextual boundaries, the likelihood of activities of interest, the relevant contextual events, and the recognition schemes available.

4.2 Trade-offs with explicit and implicit interaction

Recognizing activities of daily living can be based on data acquired through explicit or implicit interaction with the user. The decision on which approach to pursue is based on the classification and recognition accuracy of the corresponding technique, and on the resource constraints of the feature extraction and selection components for an optimal deployment. For example, sampling at 100Hz on a triaxial accelerometer and transmitting the raw data to a gateway base station for further processing will incur minimal computational overhead, but will be very expensive from a communication point of view (about 10MB per hour). By carrying out some of the data processing and analysis on the sensor, the amount of communication will be reduced. See the two deployment scenarios in Figure 2. Obviously, from a power consumption perspective there are various trade-offs to be investigated for an optimal deployment.

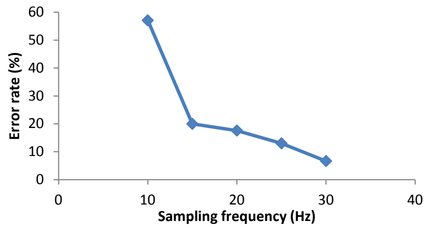

(a) Error trade-off

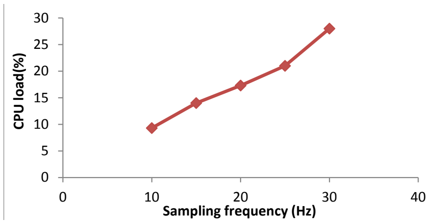

(b) Performance trade-off

Fig. 3: Sampling rate vs. recognition (a) and performance (b) on the SunSPOT

5 Experimental evaluation

We implemented two use cases: (1) fall detection and (2) step counting, and designed the activity recognition building blocks using a modular component based approach to simplify their distributed deployment. Each of these components has been profiled on various platforms. In our experiments, we used the SunSPOT sensor and an HTC Android mobile phone for profiling. We analyzed for each component on each platform the trade-offs of sampling frequency against:

1. Recognition accuracy
2. Computational complexity
3. Communication overhead and latency
4. Power consumption
5. Size and memory consumption

Similar analyses were carried out for trade-offs against the size of the sliding window, etc, but due to space considerations, we only provide in Figure 3 the results of the first two trade-off analyses for the step-counting components (features and feature classifiers) all running on the SunSPOT sensor. The figure shows both trade-offs of interest, i.e. (1) *recognition rate*, to compare different algorithms and configurations (e.g. size of sliding window and use of certain filters), and (2) *performance impact*, to decide which components to deploy and on which platform (computation vs. communication trade-off). Other trade-offs − also not shown here − investigate scenarios with all the processing done on a gateway and intermediate deployments to compare the network overhead and power consumption vs. the sampling frequency. These kind of trade-offs help us to find Pareto optimal deployments and configurations for activity recognition.

6 Discussion and future work

In our work, we are not necessarily aiming to improve the recognition rate for certain kinds of behavior and activities with complex algorithms. Rather, we are interested in finding the trade-offs between different human activity recognition components for feature selection, extraction and classification and (1) their recognition rate and (2) their resource impact for distributed deployments. We briefly discussed our *Situation Studio* tool support to model activity transitions and contextual background allowing us link that with possible recognition techniques. The techniques are implemented as modular software building blocks which can be dynamically configured, composed and deployed on our component based middleware platform that runs on sensors, smartphones and backend systems. The effects of deploying these components are profiled on each of these platforms, which helps us to find trade-offs for a distributed deployment of these component considering both recognition accuracy as well as the performance impact. As future work, we will investigate metrics for analyzing the influence of contextual background knowledge, the non-intrusiveness with explicit vs. implicit interaction.

References

1. L. Atallah, B. Lo, R. King, and Guang-Zhong Yang. Sensor positioning for activity recognition using wearable accelerometers. *Biomedical Circuits and Systems, IEEE Transactions on*, 5(4):320 –329, aug. 2011.
2. A. Avci, S. Bosch, M. Marin-Perianu, R. Marin-Perianu, and P. Havinga. Activity recognition using inertial sensing for healthcare, wellbeing and sports applications: A survey. In *ARCS Workshops*, pages 167–176, 2010.
3. A. Aztiria, A. Izaguirre, and J. Augusto. Learning Patterns in Ambient Intelligence Environments: A Survey. *Artificial Intelligence Review*, 34(1):35–51, June 2010.
4. L. Chen and C. Nugent. Ontofarm: An ontology-based framework for activity recognition and model evolution. *ERCIM News*, 2011(87), 2011.
5. M. Kangas, A. Konttila, P. Lindgren, I. Winblad, and T. Jamsa. Comparison of low-complexity fall detection algorithms for body attached accelerometers. *Gait & Posture*, 28(2):285–291, August 2008.
6. M. Keally, G. Zhou, G. Xing, J. Wu, and A. Pyles. Pbn: towards practical activity recognition using smartphone-based body sensor networks. In *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems*, SenSys '11, pages 246–259, New York, NY, USA, 2011. ACM.
7. R. Kelley, C. King, A. Ambardekar, M.N. Nicolescu, M. Nicolescu, and A. Tavakkoli. Integrating context into intent recognition systems. In *ICINCO (2)*, pages 315–320. INSTICC Press, 2010.
8. E. Kim, S. Helal, and D. Cook. Human activity recognition and pattern discovery. *IEEE Pervasive Computing*, 9(1):48–53, January 2010.
9. T. Miyaki, J. Czerny, D. Gordon, and M. Beigl. Energy-efficient activity recognition using prediction. *2012 16th International Symposium on Wearable Computers*, 0:29–36, 2012.
10. N. Noury, A. Fleury, P. Rumeau, A.K. Bourke, G.O. Laighin, V. Rialle, and J.E. Lundy. Fall detection - principles and methods. In *Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE*, pages 1663 –1666, aug. 2007.
11. L. Wang, T. Gu, H. Chen, X. Tao, and J. Lu. Real-time activity recognition in wireless body sensor networks: From simple gestures to complex activities. In *Embedded and Real-Time Computing Systems and Applications (RTCSA), 2010 IEEE 16th International Conference on*, pages 43 –52, aug. 2010.

12. C. Wren and E. Tapia. Toward scalable activity recognition for sensor networks. In *LoCA*, pages 168–185, 2006.
13. Z. Yan, V. Subbaraju, D. Chakraborty, A. Misra, and K. Aberer. Energy-Efficient Continuous Activity Recognition on Mobile Phones: An Activity-Adaptive Approach. In *16th International Symposium on Wearable Computers (ISWC)*, 2012.
14. X. Yu. Approaches and principles of fall detection for elderly and patient. In *e-health Networking, Applications and Services, 2008. HealthCom 2008. 10th International Conference on*, pages 42 –47, july 2008.
15. P. Zappi, C. Lombriser, T. Stiefmeier, E. Farella, D. Roggen, L. Benini, and G. Tröster. Activity recognition from on-body sensors: accuracy-power trade-off by dynamic sensor selection. In *Proceedings of the 5th European conference on Wireless sensor networks*, EWSN'08, pages 17–33, Berlin, Heidelberg, 2008.