

OPPORTUNITY: Towards opportunistic activity and context recognition systems

Daniel Roggen, Kilian Förster, Alberto Calatroni,
Thomas Holleccek, Yu Fang, Gerhard Tröster
Wearable Computing Laboratory, ETH Zürich
surname@ife.ee.ethz.ch

Paul Lukowicz, Gerald Pirkl,
David Bannach, Kai Kunze
Embedded Systems Laboratory
University of Passau
name.surname@uni-passau.de

Alois Ferscha, Clemens Holzmann,
Andreas Riener
Institute for Pervasive Computing
Johannes Kepler University Linz
surname@pervasive.jku.at

Ricardo Chavarriaga, Jose del R. Millán
Laboratory of Non-Invasive Neuroprosthetics
and Brain Computer Interaction
Ecole Polytechnique Fédérale de Lausanne
name.surname@epfl.ch

Abstract

Opportunistic sensing allows to efficiently collect information about the physical world and the persons behaving in it. This may mainstream human context and activity recognition in wearable and pervasive computing by removing requirements for a specific deployed infrastructure. In this paper we introduce the newly started European research project OPPORTUNITY within which we develop mobile opportunistic activity and context recognition systems. We outline the project's objective, the approach we follow along opportunistic sensing, data processing and interpretation, and autonomous adaptation and evolution to environmental and user changes, and we outline preliminary results.

1 Introduction

Sensor networks allow unobtrusive sensing of the physical world and of persons in it [23]. Wearable and pervasive computing interprets sensed data in terms of *context* [9] to provide smart assistance and context-aware ambient intelligence (AmI) environments. Complex human activities and gestures [8], and location are important aspects of context. Examples of activity- or context-aware applications are found in the industrial domain [25], healthcare [26], or interactive museums [13].

Context recognition consists of sensor signal acquisition, and its classification into a set of output classes (context) with machine learning techniques¹[27]. The prevailing approach is application-specific sensor deployment. Thus the

¹This step includes pre-processing, data segmentation into relevant sections, and feature extraction to reduce the dimensionality of the sections.

sensor signal to context mapping is known at design time. For a widespread use of context-aware systems, application specific deployment is not desirable. Users are at times in highly instrumented environments and at other times in places with little sensor infrastructure. Also, users carry a more or less random set of sensor enabled devices, such as mobile phones (e.g. with GPS and motion sensors), watches (also including various sensors), headsets, or intelligent garments (e.g. shoe worn motion sensors are commercially available). As the user leaves devices behind, picks up new ones and changes his outfit, the sensing environment changes. The on-body location of the sensors also varies, with a mobile phone at times in a pocket, in a backpack, or held in the hand. Finally sensors can fail and AmI environments will undergo upgrades over time, as new sensor-enabled devices (e.g. surveillance sensors, smart-floors) are introduced. Thus, sensing is better seen as *opportunistic*.

We envision *opportunistic activity recognition systems* that use opportunistic sensing in order to infer the user's activities and context. The key challenges are to gear opportunistic sensing towards activity and context recognition, to interpret data without assuming a-priori known sensor sets, and to a larger extent without assuming that sensor signal to context mapping is known at design time.

We investigate this within the newly started European FP7 FET-Open research project OPPORTUNITY [1] (February 2009-February 2011). In this paper we describe the objectives and key approach of the project (section 2). We detail a few initial results along opportunistic sensor processing (section 3), and autonomous system evolution (section 4).

2 Opportunistic activity/context recognition

OPPORTUNITY is a 3-year long EU FP7 project under FET-Open funding with four partners collaborating to ...develop mobile systems to recognize human activity and user context with dynamically varying sensor setups, using goal oriented, cooperative sensing. We refer to such systems as opportunistic, since they take advantage of sensing modalities that just happen to be available, rather than forcing the user to deploy specific, application dependent sensor systems.

Two major aspects are jointly investigated:

Opportunistic sensing: The system must take advantage of the available sensing infrastructure, be that ambient sensors or sensors included in on-body devices. The sensors must be coordinated and must self-organize into goal oriented sensor assemblies to achieve efficient acquisition of information relevant to the activities or contexts to recognize.

Data interpretation: The sensor signals must be interpreted in terms of activities, gestures or location. In a static system machine learning techniques learn the mapping between signal patterns and context at design time. This is not possible with opportunistic sensing. Thus context recognition must be immune to a number of variations such as sensor placement and orientation. It must also tolerate, and take advantage when possible, of changes in the sensing environment, such as when new sensors are discovered as part of an infrastructure upgrade, or sensors are removed.

2.1 The OPPORTUNITY approach

On a high level, we envision OPPORTUNITY as an interplay between:

1. *Opportunistic sensing* collecting information about the physical world in which the user behaves.
2. The instantiation of a *parameterized context recognition chain* based on the available sensors and their characteristics.
3. System *adaptation* during operation to sensing environment changes by reconfiguring the activity recognition chain to achieve the desired recognition goal.

In figure 1 we outline the OPPORTUNITY system. It comprises 6 core ideas and principles, detailed below, that contribute to opportunistic context awareness.

2.1.1 Abstract recognition goal

The activity or context recognition goal is formulated in an abstract manner (fig. 1 top). This is in contrast to current systems where the recognition task is hard-coded in the system. The goal formulation says *what* should be recognized, but does not specify *how*. This allows the system to autonomously configure itself to use the appropriate and available resources.

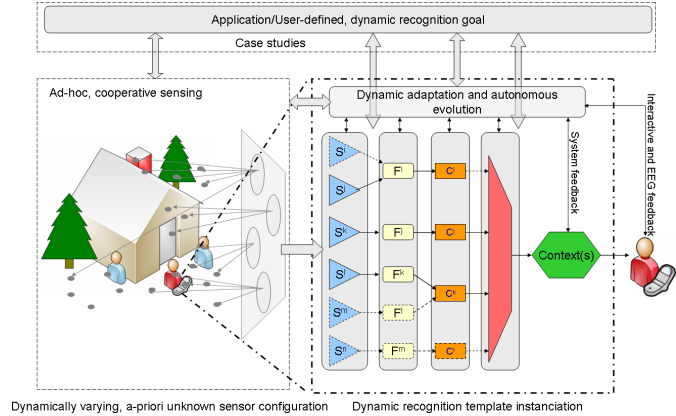


Figure 1. The overall OPPORTUNITY system

For example, in order to prevent dangerous manipulations, a recognition goal may be to detect grasping manipulative activities for anybody located in a warehouse with dangerous goods, with highest accuracy, using any ambient or wearable sensors. In a lifestyle application that monitors physical exercise the recognition goal would relax the constraints on accuracy and favour less obtrusive, mostly wearable sensors.

The actual methodologies to formulate abstract recognition goals are investigated in the project with emphasis on goal representations, and methods to convert the goal into a sensing mission [12].

2.1.2 Opportunistic sensing

The recognition goal is translated into a coordinated sensing mission (fig. 1 left). A number of approaches exist for coordinated emergence of sensing networks [2, 15, 3, 10] and autonomous composition [14, 24] and evolution [21, 5, 20] of services. **Autonomous self-organization** according to the sensing mission results in a goal-oriented sensing ensemble providing the relevant information to the mobile system from wearable or ambient sensors. As an example, if the context to recognize are manipulative activities, the information will come only from sensors capable of providing relevant information. This means all body-worn sensors capable of sensing motion, such as inertial sensors, ultrasound or ultrawideband tracking from on-body tags, or muscle activity sensors (EMG). Sensors such as temperature or presence sensors would be excluded as they do not provide relevant information to infer manipulative gestures. **Self-description and self-characterization** underly self-organization. Sensor nodes must advertise their characteristics, such as the physical quantity they sense (e.g. sound, motion), with which parameters (e.g. accuracy, sample rate), exact location or at least gross location (e.g. body-worn sensor v.s. ambient sensors). Some descriptive parameters must be inferred while the system operates by self-characterization. For example, a cell-phone can be carried on body, or left lying on a table. However it can self-

characterize its location by recognizing typical movement patterns and advertise this information (see section 3).

The main research efforts address the components for an infrastructure-free opportunistic sensing system, optimized for human activity and context recognition. This includes sensor markup, ensemble coordination, and self-description and self-characterization methods [12].

2.1.3 Signal conditioning and feature abstraction

A typical context recognition chain consists of pre-processing, feature extraction, feature classification, and decision fusion (fig. 1 right). In order to allow variability in the sensing infrastructure we investigate signal conditioning and sensor independent features (see section 3 for illustrative results).

Signal conditioning consists in re-defining commonly used features in a way as to make them less sensitive to variations. A simple example is the acceleration magnitude captured by a 3-axis accelerometer that is independent of the sensor orientation. Another example is to combine sensors so that one compensates the effect of variations on the other.

Sensor independent features. Different physical quantities may provide related information about an activity. For example, in order to classify hand motion into gestures the hand trajectory is required. Several sensors can be used: e.g. on-body inertial sensors, or visual tracking. Thus, hand trajectory is a sensor independent feature that can be obtained from different sources. This relaxes the constraints on sensor availability for a desired recognition goal. We envision this method in the form of a table that list for a number of the most common recognition problems groups of sensor types that can be combined to obtain the abstract features, as well as the corresponding transformation algorithm. Thus, according to the recognition problem and sensor availability, this table allows to find which sensors to combine to obtain the required abstract features.

2.1.4 Opportunistic machine learning methods

Signal conditioning and abstract features cannot filter all variability. Thus, classifiers and decision fusion must ensure graceful degradation in case of variability, and have the ability to dynamically exploit additional and/or improved sources according to their information content. In OPPORTUNITY we envision ensemble classifiers [22] with dynamic combination of classifiers at run-time [29] for fault-tolerant and flexible sensor fusion. We will investigate the benefits brought about by sensor self-description and self-characterization. We present some initial results in dynamic ensembles for opportunistic activity recognition in section 4.1.

2.1.5 Autonomous evolution and adaptation

The sensing infrastructure of AmI environments changes over time as sensors fail or the infrastructure is upgraded with new sensors. New sensors may provide information

relevant to some of the contexts. However the system needs to make sense of this information (e.g. if the new sensor is a sensor without or with incomplete self-description capabilities). Slow, long-term changes in signal characteristics (e.g. sensors affected by harsh environments) and user activity patterns (e.g. due to change of habits, ageing) also occur. This translates into changing mappings between sensor signals and context. These changes are not predictable, but an autonomous opportunistic system must be immune to them at least, or capable of taking advantages of these changes at best.

In order to adapt to changing sensor numbers, we can exploit the fact that when two sensors show correlated signals it is likely they are measuring the same physical quantity at the same location. The longer the correlation period, the higher is the likelihood [19]. This allows us to devise autonomously and continuously, at run-time, a *substitution table*. It indicates the degree to which two sensors are believed to be equivalent. It allows to replace a sensor that fails by one or more other sensors that are considered as equivalent. It also allows to learn how to make use of newly discovered sensors that are correlated with existing ones. Correlations can also be assessed in the feature space, and transformations between signals or features can be included. On a higher level, the existing system can also provide estimated ground truth labels to control the training of classifiers operating on the signals provided by newly discovered sensors (see section 4).

Slow changes in the mapping between sensor signals and context translates by a drift of the points corresponding to activity classes in the feature space (concept drift [28]). By monitoring drift over time, the classifier parameters may be adjusted accordingly. Adaptation can be done after a calibration routine, or by taking advantage of many repetitive occurrences of context in daily life (see section 4).

2.1.6 Opportunistic user feedback

A user can report errors in the system's context-aware behavior. An opportunistic system can take advantage of this occasional feedback to collect error statistics and adjust its recognition behavior (fig. 1 far right). A mobile device allows interactive feedback. However, more promising is the use of human generated signals related to cognitive states to provide learning signals to artificial systems. In particular, the detection of error-related EEG correlates (brain signal patterns occurring when a system deviates from expected behavior) is a promising approach to derive an endogenous measure of system performance to guide system adaptation [7, 11]. By combining this with semi-supervised learning methods [6], classifiers can be efficiently trained from a large amount of unlabelled data and a small proportion of opportunistically labelled data. This may allow OPPORTUNITY to turn into an autonomic system capable of self-improvement.

2.2 Validation

We will apply opportunistic principles for simple primitives of context and activity at first, such as presence, location, modes of locomotion, and hand gestures. We will then combine primitives into composite activities, such as object manipulation or simple interactions between humans. By combining these building blocks into larger scenarios we will investigate the extent up to which these approaches scale in more realistic application scenarios, such as indoor activity monitoring, or lifestyle monitoring. One particular test-bed is the development of robust adaptive Brain-Computer Interface (BCI) systems, as an example of complex cognitive context recognition. EEG-based BCI is highly sensitive to noise and electrode contact with skin. Applying OPPORTUNITY methods to select appropriate electrode channels according to the current situation will tell us about the generalizability of the methods to other problem domains.

3 Signal conditioning and feature abstraction

In ongoing work towards activity recognition from sensor sets opportunistically discovered on-body (e.g. phones, PDAs, watches, headsets), we investigated signal conditioning and sensor abstraction building blocks that can be combined with other methods outlined in this paper. In particular, an opportunistic recognition system must address the problem of placement and orientation independence. Thus one has to be able to detect on-body location if the device is worn on different body parts, one has to deal with displacement issues, and if the device is not worn it is useful to know its symbolic location.

3.1 On-body location

Wearable devices can be in a number of different locations (e.g. headset on the head or in one of many pockets). Previously we have shown on a small data set how to recognize a set of on-body locations using an accelerometer signal during walking and arbitrary activities [17]. We verified on-body location recognition accuracy on a dataset containing 9 hours from real life activity by three users ranging from a 70 y.o. housewife to a 28 y.o. male student. The maximum accuracy achieved for 6 min windows is 82 % for 5 locations (head, wrist, torso, front pocket, back pocket) (fig. 2).

3.2 On-Segment Displacement

We presented a set of heuristics that significantly increase the robustness of motion sensor-based activity recognition with respect to sensor displacement within a single body part (e.g. lower arm). Within certain limits and with modest quality degradation, motion sensor based activity recognition can be implemented in a displacement tolerant way by taking into account physical principles of body mechanics [18]. After evaluation on a set of synthetic lower

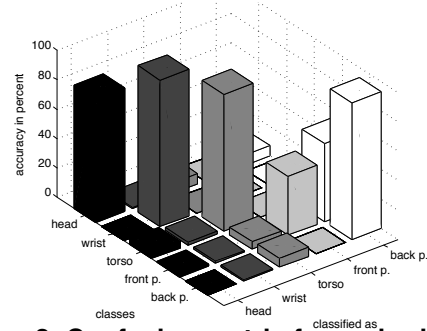


Figure 2. Confusion matrix for on-body location detection using hidden Markov models.



Figure 3. Setup for displacement robust recognition.

arm motions we extended this approach to modes of locomotion problem (sensors on the upper leg) and to a set of exercises performed on various gym machines (sensors placed on the lower arm, see fig 3). In this example our heuristic raises the displaced recognition rate from 24% for a displaced accelerometer, which had 96% recognition when not displaced, to 82%.

3.3 Symbolic Object Location

In [16] we describe a novel method for symbolic location discovery of simple objects. It requires no infrastructure. It relies on sensors typical in ambient intelligence environments and smart objects (acceleration, sound). Objects emit short, narrow frequency “beeps”. Sound reflection in the environment is typical to the location and recognized by audio fingerprinting (fig. 4). We assessed this method on specific locations such as “on the couch”, “in the desk drawer”, “inner jacket pocket”, “outer jacket pocket”, and on abstract locations such as “in closed wood compartment”, “on open iron surface”. We did a study with over 1200 measurements, 35 specific locations in 3 rooms, and 12 abstract location classes. On the 16 locations in the largest room the recognition rate is 96 %. On the whole 35 locations it is 81 %, however the correct location is in the 3 top picks of the system 94 % of the times.

4 Autonomous evolution and adaptation

Autonomous adaptation to changes in the sensing environment is desired without (or at least with minimal) inter-




Figure 4. Audio fingerprint of a beeping object placed on a carpet (left) or desk (right).




Figure 5. Left: performance according to ensemble size. Right: dynamic ensemble adaptation at runtime (black=active sensors) allows to reach a target performance despite sensors failures. Time is from top to down.

vention. Thus the system evolves according to the sensing situation and is capable of operating in open-ended environments. In ongoing work we investigate: ensemble adaptation when sensors are lost; exploitation of newly discovered sensors (e.g. for sensors insufficiently capable of self-description or “legacy” sensors); and autonomous classifier adaptation to cope with slow, long-term changes in sensor and activity properties (concept drift) e.g. due to sensor ageing or changes in the way users perform activities (e.g. proficiency increase, ageing).

4.1 Ensemble modelling and dynamic adaptation

We have previously investigated empirical performance models of ensemble classifiers collaborating to recognize activities. Up to 60 sensors collaborated to recognize a set of 10 manipulative gestures in an automotive environment [29]. We derived empirical performance models suited for online use and showed that classifier ensembles provide intrinsic robustness to noise and faults, and that performance could be scaled in a simple case with the number of sensors (fig. 5 left). This performance model allows dynamic adaptation of the ensemble to the current sensor availability, while guaranteeing a desired classification accuracy (fig. 5 right). Opportunistic communication aspects now need to be factored in [4].

4.2 Exploitation of newly discovered sensors

We investigate the use of the current system to detect context occurrences and use this as ground-truth to train a




Figure 6. Performance v.s. sensor turnover: new sensors are trained by the system and replace one sensor of the ensemble at each iteration.

classifier operating on the signals of the new sensors. We analyzed a worst-case on the same scenario as above. A sensor is added to the system. It is trained by the current sensors in the ensemble upon context occurrences. It then replaces one sensor in the ensemble to keep its size equal. This is repeated with further additional sensors. Figure 6 shows the performance of the system according to the number of replaced sensors. The performance decreases rapidly when the ensemble contains one sensor (30% of original performance after 300 iterations). However with as few as three sensors, the ensemble maintains high accuracy (98% of original performance after 300 iterations). This approach can be combined with occasional accurate user-generated ground truth. Thus, we investigate the ratio of ground truth labels that may balance this performance drop.

4.3 Autonomous classifier adaptation

Changing sensor signal to context mappings affect typical machine learning methods. We investigate autonomous classifier adaptation by which, upon detection of context occurrence, the classifiers are re-trained to better model the corresponding feature point. In the case of a Nearest Class Center (NCC) classifier, the class center in the feature space that corresponds to the newly recognized event is displaced towards that point. In figure 7 we show the probability of reaching a preset accuracy versus the accuracy of the classifier before the unsupervised adaptation process for a two class problem. With only a slightly better initial true positive rate than chance (51%) the probability of improving the classifier is already 87%. The classifiers that were successfully adapted reached more than 95% accuracy on average. The adaptation dynamics of three different class centers from a NCC classifier to two dimensional Gaussian distributions are shown in figure 7. Class centers are attracted by high density regions in the distributions.

5 Conclusion

Opportunistic sensing is a promising paradigm to collect information on a large scale about events occurring in the physical world, and thus particularly interesting for ambient intelligence environment and wearable systems where the nature and availability of sensor-enabled devices often

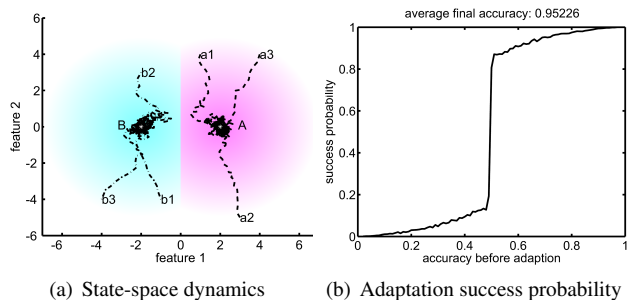


Figure 7. Classifier adaptation

cannot be controlled. In these application domains sensor information must be interpreted in terms of human context or activity. In this paper we outlined the newly started EU project OPPORTUNITY within which we develop methods for opportunistic human activity and context recognition. We presented results of a few methods we investigate to take advantage of opportunistic sensing and bring it towards opportunistic recognition of human contexts and activities. These results outline our initial investigations but do not represent the final OPPORTUNITY system as the project just recently started. We invite all readers to access our website [1], where we will publish project updates.

Acknowledgment

The project OPPORTUNITY acknowledges the financial support of the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET-Open grant number: 225938.

References

- [1] Opportunity project web site, February 2009. <http://www.opportunity-project.eu/>.
- [2] J. Beal and J. Bachrach. Infrastructure for engineered emergence on sensor/actuator networks. *IEEE Intelligent Systems*, 21(2):10–19, 2006.
- [3] N. Biccocchi, M. Baumgarten, M. Mulvenna, R. Kusber, and F. Zambonelli. Self-organizing knowledge networks for autonomic communication services. In *IEEE Conference on Systems, Man, and Cybernetics*, 2007.
- [4] C. Boldrini, C. Marco, and A. Passarella. Modelling data dissemination in opportunistic networks. In *CHANTS '08: Proceedings of the third ACM workshop on Challenged networks*, pages 89–96, New York, NY, USA, 2008. ACM.
- [5] I. Carreras, I. Chlamtac, F. De Pellegrini, and D. Miorandi. Bionets: Bio-inspired networking for pervasive communication environments. *IEEE Trans. Veh. Tech.*, 56(1):218–229, 2007.
- [6] O. Chapelle, B. Schölkopf, and A. Zien. *Semi-Supervised Learning*. MIT Press, Cambridge, MA, 2006.
- [7] R. Chavarriaga, P. W. Ferrez, and J. del R. Millán. To err is human: Learning from error potentials in brain-computer interfaces. In *International Conference on Cognitive Neurodynamics*, 2007.
- [8] N. Davies, D. P. Siewiorek, and R. Sukthankar. Special issue: Activity-based computing. *IEEE Pervasive Computing*, 7(2):20–21, 2008.
- [9] A. K. Dey. Understanding and using context. *Personal and Ubiquitous Computing Journal*, 5(1):4–7, 2001.
- [10] F. Dressler. *Self-organization in sensor and actor networks*. Wiley, 2007.
- [11] P. W. Ferrez and J. del R. Millán. Error-related eeg potentials generated during simulated brain-computer interaction. *IEEE Trans Biomed Eng.*, 55:923–929, 2008.
- [12] A. Ferscha, A. Riener, M. Hechinger, R. Mayrhofer, M. dos Santos Rocha, A. Zeidler, and M. Franz. Peer-it: Stick-on solutions for networks of things. *Pervasive and Mobile Computing Journal*, 2008.
- [13] M. Fleck, M. Frid, T. Kindberg, E. O’Brien-Strain, R. Rajani, and M. Spasojevic. From informing to remembering: Ubiquitous systems in interactive museums. *IEEE Pervasive Computing*, 1(2):13–21, 2002.
- [14] K. Fujii and T. Suda. Semantics-based dynamic web service composition. *International Journal of Cooperative Information Systems*, 15(3):293–324, 2006.
- [15] R. Haesevoets, B. Van Eylen, D. Weyns, A. Helleboogh, T. Holvoet, and W. Joosen. Coordinated monitoring of traffic jams with context-driven dynamic organizations. In *Int. Conf. on Engineering-Mediated Multiagent Systems*, 2007.
- [16] K. Kunze and P. Lukowicz. Symbolic object localization through active sampling of acceleration and sound signatures. *Ubicomp*, 2007.
- [17] K. Kunze and P. Lukowicz. Using acceleration signatures from everyday activities for on-body device location. *Wearable Computers, 2007 11th IEEE International Symposium on*, pages 115 – 116, Sep 2007.
- [18] K. Kunze and P. Lukowicz. Dealing with sensor displacement in motion-based onbody activity recognition systems. *Proc. 10th Int. Conf. on Ubiquitous computing*, Sep 2008.
- [19] J. Lester, B. Hannaford, and G. Borriello. ” Are You with Me?”-Using Accelerometers to Determine If Two Devices Are Carried by the Same Person. *LECTURE NOTES IN COMPUTER SCIENCE*, pages 33–50, 2004.
- [20] D. Miorandi, L. Yamamoto, and P. Dini. Service evolution in bio-inspired communication systems. *Int. Trans. Syst. Sc. and Appl.*, 2(1):51–60, 2006.
- [21] T. Nakano and T. Suda. Applying biological principles to designs of network services. *Applied Soft Computing*, 3:870–878, 2007.
- [22] R. Polikar. Ensemble based systems in decision making. *IEEE Circuits and Systems Magazine*, 6(3):21–45, 2006.
- [23] D. Puccinelli and M. Haenggi. Wireless sensor networks: applications and challenges of ubiquitous sensing. *IEEE Circuits and Systems Magazine*, 5(3):19–31, 2005.
- [24] R. Quitadamo, F. Zambonelli, and G. Cabri. The service ecosystem: Dynamic self-aggregation of pervasive communication services. In *1st ICSE Workshop on Software Engineering of Pervasive Computing Applications, Systems and Environments (SEPCASE'07)*, 2007.
- [25] T. Stiefmeier, D. Roggen, G. Ogris, P. Lukowicz, and G. Tröster. Wearable activity tracking in car manufacturing. *IEEE Pervasive Computing*, 7(2):42–50, 2008.
- [26] M. Tentori and J. Favela. Activity-aware computing for healthcare. *IEEE Pervasive Computing*, 7(2):51–57, 2008.
- [27] J. Ward, P. Lukowicz, G. Tröster, and T. Starner. Activity recognition of assembly tasks using body-worn microphones and accelerometers. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 28(10):1553–1567, 2006.
- [28] G. Widmer and M. Kubat. Learning in the presence of concept drift and hidden contexts. *Mach. Learn.*, 23(1):69–101, 1996.
- [29] P. Zappi, C. Lombriser, E. Farella, D. Roggen, L. Benini, and G. Tröster. Activity recognition from on-body sensors: accuracy-power trade-off by dynamic sensor selection. In R. Verdone, editor, *5th European Conf. on Wireless Sensor Networks (EWSN 2008)*, pages 17–33. Springer, 2008.