
Robust activity recognition for assistive technologies: Benchmarking machine learning techniques

**Ricardo Chavarriaga, Hesam Sagha,
Hamidreza Bayati, José del R. Millán,**
Non-Invasive Brain-Computer Interface (CNBI)
Ecole Polytechnique Fédérale de Lausanne
ricardo.chavarriaga@epfl.ch

**Daniel Roggen, Kilian Förster,
Alberto Calatroni, Gerhard Tröster**
IFE, Wearable Computing Lab
ETH Zürich
daniel.roggen@ife.ee.ethz.ch

Paul Lukowicz, David Bannach
Embedded Systems Laboratory
University of Passau
paul.lukowicz@uni-passau.de

Marc Kurz, Gerold Hölzl, Alois Ferscha
Institute for Pervasive Computing
Johanes Kepler University, Linz
ferscha@soft.uni-linz.ac.at

Abstract

An increasing need for healthcare provision and assistive technologies (AT) calls for the development of machine learning techniques able to cope with the variability inherent to real-world deployments. In the particular case of activity recognition applications sensor networks may be prone to changes at different levels ranging from sensor data variability to network reconfiguration. Robust methods are required to deal with those changes providing graceful degradation upon failure or self-configuration and adaptation capabilities that ensure their proper operation for long periods of time. Currently there is a lack of common tools and datasets that allow for replicable and fair comparison of different recognition approaches. We introduce a large database of human daily activities recorded in a sensor-rich environment. The database provides large amount of instances of the recorded activities using a significant number of sensors. In addition, we reviewed some of the techniques that have been proposed to cope with changes in the system, including missing data, sensor location/orientation change, as well as the possibility to exploit data from unknown discovered sensors. These techniques have been tested in the aforementioned datasets showing its suitability to emulate different sensor network configurations and recognition goals.

1 Introduction

There is an increasing need for healthcare provision and assistive technologies (AT). This is mainly driven by the aging population in developed countries, as well as enabling factors such as technological advances in sensing and portable computing devices and wireless communication. These AT devices can be used to support people with motor or cognitive disabilities in the form of assisted mobility or prosthesis [1], hearing aids, among others. Another main line of AT devices concerns human activity recognition, these include monitoring of people's activities for health assessment and rehabilitation; as well as intervention tools to promote healthy daily routines. Another example are wearable health assistants for Parkinson's disease patients in order to prevent *freeze of gait* [2]. Alternatively, activity recognition can also be applied in healthcare applications to provide context-aware information to clinicians and nurses in hospitals [3].

In this paper we focus on activity recognition for AT applications, noticing that they pose important challenges to the machine learning community. Indeed, real-world deployment of these systems

suffers from existing limitations of the current approaches undertaken by the community. Indeed, recognition systems are typically designed and tested for narrowly defined tasks, assuming fixed and well-known sensor configurations[4]. However, these assumptions cannot be enforced in practical applications where the system is expected to operate for long periods of time, or sensor placement has to be done in a daily basis, sometimes by the user herself. Moreover, sensor-enabled devices such as PDAs or smartphones, for which a continuous change of location is expected, can be included in the network [5]. Similarly, as users move around different environments, ambient sensors can become available or disappear with time (e.g. moving from a highly instrumented environment such as a “smart home” to the street where no ambient sensors are present). Finally, long-term system operation is also prone to sensor and communication failure, as well as the possibility of upgrading when new sensor technology is available. All this requires the system’s ability to cope with those changes in the least-intrusive way to avoid the cost of re-training and redeployment. Machine learning (ML) techniques for AT should then be able to deal with these different sources of variability, and have the capacity to adapt to changes in the sensor network as well as changes in the person behaviour (e.g. as a patients go through rehabilitation therapy her motion patterns may change) in a transparent, reliable manner.

All these factors calls for techniques robust to sensor signal degradation, parameters and network configuration changes by means of self-adaptation and re-configuration, ideally in an unsupervised manner. Moreover, proper assessment of newly developed recognition techniques requires the availability of benchmarking data that allows replicability of the testing procedures for different approaches, while capturing the variability that characterise real-world activity recognition tasks. In this paper we highlight the lack of such datasets and introduce a new dataset of human daily activities in a sensor-rich environment specifically conceived to address these issues (Section 2). Furthermore, we discuss some of the ML approaches that have been proposed to provide the desired adaptation capabilities, tested on several configurations of this dataset (Section 3).

2 Benchmarking activity recognition: The Opportunity dataset

Unlike other applications of machine learning there are no established benchmarking problems for activity recognition. Typically, each research group tests and reports the performance of their algorithms on their own datasets using experimental setups specially conceived for that specific purpose. For this reason, it is difficult to compare the performance of different methods or to assess how a particular technique will perform if the experimental conditions change (e.g. in case of sensor failure or changes in sensor location). In contrast, for other application fields often exist publicly available datasets that allow the comparison of different machine learning algorithms on the very same conditions. This comparison is often done in the form of competitions and challenges providing a fair comparison of the proposed techniques. These competitions have taken place using datasets from diverse fields including computer vision, bioinformatics, computational neuroscience or brain-computer interfaces.

Some available datasets for activity recognition do exist, but they are usually specific to an activity recognition purpose. Among the existing datasets we can mention the *PlaceLab dataset*, probably the most popular dataset in the pervasive/ubiquitous computing field. It focuses on ambient and object sensing collected from subjects behaving for up to a week in an environment with multimodal sensors [6]. It contains long recordings but it doesn’t include a high number of activity instances. *Van Kasteren’s dataset* [7] features longer recordings (month-long) but has fewer sensors. It uses digital or binary sensors like reed switches to record interactions with furniture or objects of interest. No information about modes of locomotion or posture is available. The *Darmstadt routine dataset* used for unsupervised activity pattern discovery [8], is a long recording from body activity collected by the Porcupine system [9]. The *TUM Kitchen dataset* focuses on video-based activity recognition [10], and also contains RFID and reed switch data, but it does not include on-body sensors.

The lack of more general databases can be explained by the difficulty to conceive and record a dataset that is able to reflect the complexity and variability of situations that may be encountered in healthcare scenarios. Moreover, proper comparison of machine learning techniques requires these datasets to provide a reasonable amount of instances for the different recorded actions and to include several subjects in order to allow the assessment of inter-subject variability. In addition, if the

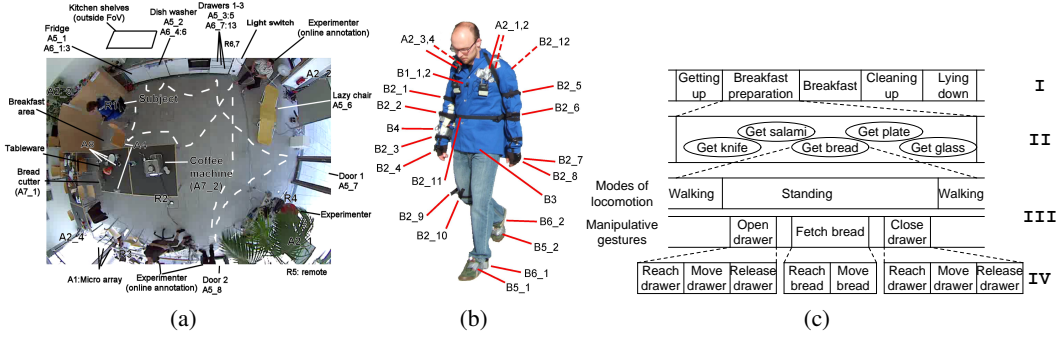


Figure 1: Opportunity dataset setup. (a) View of the room from top. Dashed line: typical user trajectory in the drill run. (b) Location of the body-worn sensors on the subject. (c) Decomposition of recorded activities. Level I: High level activities. Level II decomposes high level activities. Level III: modes of locomotion and manipulative gestures. Level IV encapsulates the atomic activities forming the manipulative gestures.

database wants to be used for emulating changes in the sensor network, activities should be recorded by a large and diverse set of sensors.

In order to cope with the limitations of existing datasets we recorded naturalistic human activities in a sensor rich environment: a room simulating a studio flat with kitchen, deckchair, and outdoor access where subjects performed daily morning activities [4, 11]. We deploy 15 networked sensor systems of different origins (proprietary and custom, from different manufacturers or universities). It comprises 72 sensors of 10 modalities, integrated in the environment, in objects, and on the body¹. The setup including both ambient and on-body sensors is illustrated in Figures 1(a) and 1(b). Our aim is to provide the research community a publicly available dataset of complex, interleaved and hierarchical naturalistic activities, with a particularly large number of atomic activities (around 30'000), collected in a very rich sensor environment. The dataset contains 25 hours of recorded data from 12 subjects. The cumulative length of all the annotations on all tracks represents approximately 57 hours of labels, including posture/locomotion labels which are always present and hand interactions that often occur in overlapping fashion and including multiple objects. As an illustration, Table 1 shows the number of instances for locomotion modes and hand interactions in 19 annotated runs (from a total of 72 recorded runs). The number of available sensors and activity instances makes this dataset well suited to benchmark various activity recognition approaches, and to investigate e.g. multimodal data fusion, reasoning, or activity and scenario modelling.

For the design of the database scenario, we focus on maximising the number of activity instances that were collected, while keeping their execution naturalistic. We achieved this by relying on a high-level script and leaving free interpretation to the users, and even encouraging them to perform as naturally as possible with all the variations they were used to. During the recordings, each subject performed 5 times a run with activities of daily living (*ADL*) and one *drill* run. The *ADL* run consists of temporally unfolding situations. In each situation (e.g. *preparing sandwich*), composite activities (e.g. *cutting bread*) occur as well as atomic activities (e.g. *reach* for bread, *move* to bread cutter, *operate* bread cutter). This allows to look at activity recognition at various abstraction levels (c.f. Figure 1(c)). During the drill runs, subjects performed 20 repetitions of a predefined sequence of activities including open and close doors and drawers, turn on/off the lights or drink. This yields well-structured recordings containing a rather large number of activity instances

Given the characteristics of the database where all activities are sensed by multiple sensors of different modalities allows to emulate different sensor network configurations. For instance, since some sensors are located in close proximity to others they can be used to study robustness against sensor placement variability [12]. Alternatively, sensors of different modalities but sensing information related to a common activity allow to study the dynamic replacement of one modality by another. Finally, multiple sensors of identical modalities but from different systems allow to assess the effects of calibration, resolution, or sample rate variations. In the next section we address some of

¹A video presentation of the dataset is available at <http://vimeo.com/8704668>.

Table 1: Overall instances of modes of locomotion and hand-related activities, along with minimum, maximum, average and total duration (seconds)

	#	Min Length	Max Length	Mean Length	Total Time
Modes of locomotion					
Walk	1414	0.3	242.6	5.6	7900.5
Stand	1043	0.2	171.2	7.5	7770.9
Lie	56	0.9	166.6	21.8	1219.5
Sit	127	0.8	274.9	26.4	3349.5
Hand-environment interactions					
Ambient	3426	0.2	6.3	1.0	3313.4
Objects	3709	0.2	64.5	2.0	7399.8

the challenges for activity recognition and show how different methods have been assessed using different configurations of the presented dataset.

3 Robust ML approaches for activity recognition

As mentioned, activity recognition for real-world AT applications should be robust against possible changes in the sensor network or user behaviour. Moreover, adaptation mechanisms that avoid or reduce tedious and costly training and calibration processes are also desirable. Among the different types of variability that these systems may face we can count: *Changes at the data level*, due to loss of communication or sensor degradation; *Changes at the sensor level*, as sensors may slip or rotate during operation, be relocated by the user (e.g. an instrumented smartphone may be carried at different pockets throughout the day) or they can be placed at slightly different positions from one day to the next. Finally, *changes at the network level* can also occur as the user moves around environments with different sensors. This requires the recognition system to be able to self-configure itself upon the discovery of new sensors or the disappearance of others, and calls for tools that allow the management of the network data flow to support the dynamic routing of information and data across the network and the incorporation of training sources for enabling online training of sensor devices at runtime [13].

In the following subsections we will briefly review some of the methods that have been proposed to address some of these issues. In addition, we illustrate the use of the *Opportunity* dataset to emulate different network configurations and recognition chain setups, thus showing how it can be a useful tool for assessing the performance of newly proposed ML techniques. We present examples addressing the handling of missing data, and the automatic recognition of changes (e.g. due to sensor rotation, or displacement), as well as the automatic exploitation of newly discovered sensors.

3.1 Handling missing data

Sensor networks are prone to data loss due to disconnections, sensor failure, and transmission problems. This is particularly relevant for wearable and wireless sensors deployed in real-life scenarios. The problem of classification with missing data has been extensively studied in the fields of speech recognition and bioinformatics, but less studied in the case of activity recognition. Some methods impute data based on distance metrics [14]. In this case, a distance measure is computed between the current available sensor values and a previously stored patterns. When incoming data is missing, the closest stored pattern is imported. Therefore, the performance strongly depends on the quality and the number of patterns stored in the dataset.

We compare several techniques for handling missing data using the Opportunity dataset. We classify 16 low-level actions using data of 5 ADL sessions for one subject. Each data source is classified independently using a GMM classifier and then decisions are fused using Dempster-Shafer. Figure 2 shows the recognition performance when the sensor with missing samples is removed from the recognition chain (*Removal*), as well as several imputation techniques. *Cluster* refers to a modified version of the distance-based method proposed by [14]. The other two methods use a probabilistic approach that infers the missing samples using the conditional distribution between available and




Figure 2: Comparing techniques for handling missing data on different network configurations.

missing data sources. This inference can be performed at the raw data level or at the fusion level; RPCD and DPCD, respectively. Results show that the latter outperforms all the other methods [15].

3.2 Detecting changes in the sensor network

Another critical issue for activity recognition is detecting sensors that fail or degrade [16, 17]. This degradation or fault may be due to loss of power, change of the sensor position or due to environmental fluctuations. As an example, degradation may come from the slippage or rotation of on-body accelerometer sensors [4]. Moreover, for assistive technologies it is unrealistic to expect users to always locate the sensors at exactly the same position day after day.

Machine learning techniques can be applied to allow for detection of these anomalies improving the robustness of the recognition chain. Cumulative Sum (CUSUM) approaches detect changes in the distribution of a measured variable [18]. These methods typically monitor only one stream of signals and they decide whether it is faulty or not. However, in a network of sensors a sensor may not be faulty but its behaviour may have changed in the network. We have recently proposed an anomaly recognition technique based on a distance measure between a given sensor and the rest of the sensors in the network. Training data can be used to characterise these distances and, at operation time changes are detected when these distances are exceeded [19]. Following a classifier fusion approach, detected anomalous sensors can be removed from the fusion. We simulate changes in the network by adding rotational noise in the same setup described in the previous section. Figure 4 shows the recognition performance using acceleration sensors alone or in combination with gyroscopes; moreover we also compare Naive Bayesian fusion and Dempster-Shafer (DS) fusion techniques.

The presented approach removes the sensors detected as anomalous, other approaches cope with these changes by selecting robust features. Förster and colleagues address changes due to sensor displacement by selecting location-invariant features using genetic algorithms [20]. Finally, unsupervised learning methods can be applied to track changes in the feature space; An adaptive version of NCC classifiers has been proposed to adapt to changes in sensor location [12]. More recently, the unsupervised approach has been extended based on covariance shift estimation. It assumes that sensor changes are mainly reflected in shifts in the feature distributions and estimates this shift online using a modified version of expectation-maximisation [21].

3.3 Knowledge discovery: Exploiting unknown sensors

Activity recognition systems for assistive applications are expected to operate for long periods of time in environments susceptible of upgrades or changes. As the sensor network changes, it would be useful to be able to exploit unknown newly discovered sensors for the intended goal. We have shown that sporadic interactions with primitive sensors (typically simple ambient sensors such as reed switches) combined with behavioural assumptions can be used to confer activity recognition capabilities to a newly discovered sensor [22]. These behavioural assumptions correspond to hypotheses about likely human behaviour; for example, depending on the environment, we can assume that somebody is going to be standing while opening a window, while she will be walking towards it and away from it shortly before and after her interaction with the window.

This approach has been successfully tested in the Opportunity dataset to recognise modes of locomotion. We exploit on-body Ubisense tags and reed switches located on three drawers, a dishwasher and a fridge. These switches indicate when they are opened/closed. The behavioural assumption is

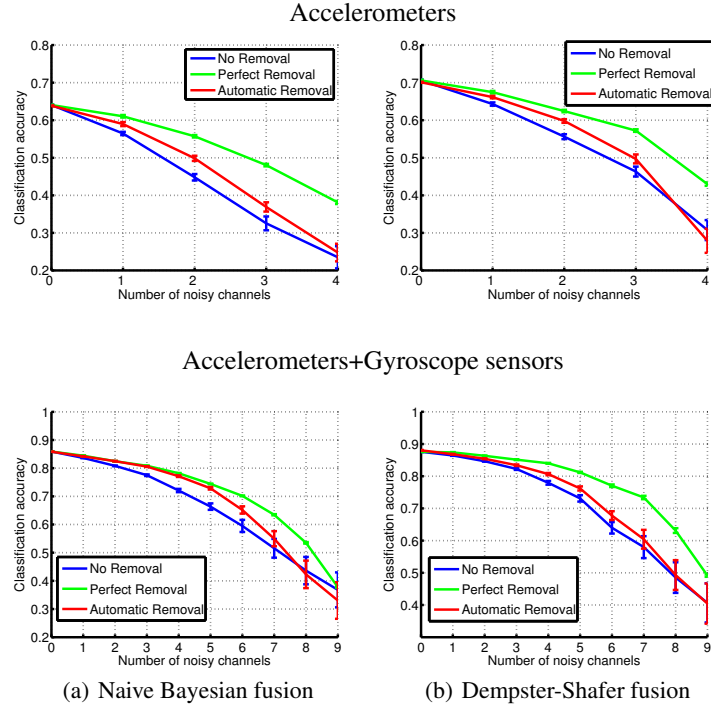


Figure 3: Change detection in sensor networks. *No Removal*, noisy sensors are included in the recognition chain. *Automatic Removal*, sensors detected as anomalous are removed from fusion. *Perfect Removal*, noisy sensors are manually removed from the fusion.




Figure 4: Exploiting unknown sensors. (a) Behavioural assumptions used to infer modes of locomotion. W=walk, S=stand, GT=ground truth, BA=behavioural assumptions. δ_S , δ_M and δ_W denote respectively the width of the generated 'Stand' label, the margin and the width of the 'Walk' label. (b) 'Stand' label purity for increasing values of δ_S . (c) Label purity for different combinations of two of the parameters, namely δ_S , the width of the 'Stand' label, and δ_M , the margin between 'Stand' and 'Walk' labels.

that when the user opens/closes the item, she stands, and that she walks shortly before and after. Figure 4(a) shows the behavioural assumptions used to infer modes of locomotion. The "purity" of the generated label, i.e. proportion of time it corresponds to ground truth is shown in Figures 4(b) and Figure 4(c). The generated labels can then be used to incrementally train an unknown sensor, placed on the user's body, so that this learns to recognize the aforementioned modes of locomotion. The challenge is to find a set of ambient sensors and robust assumptions that allow an efficient training of the unknown sensors

4 Conclusion

Machine learning techniques for AT applications face interesting challenges regarding their ability to cope with the inherent variability of real-world deployments. In the particular case of activity

recognition sensor networks may be prone to changes at different levels ranging from sensor data variability to network reconfiguration. Robust methods are then required to deal with those changes providing graceful degradation upon failure or self-configuration and adaptation capabilities that ensure their proper operation for long periods of time.

Several methods have been proposed to address these issues, but they are generally designed and tested in well-defined scenarios that assume strong constraints in terms of the network configuration. This makes difficult to infer the performance of these systems in the case of changes in the recognition chain. We claim that common tools that allow replicable and fair comparison of ML methods may constitute an important contribution towards the development of more robust approaches. For this reason, we introduce a database of human daily activities recorded in a sensor-rich environment. Besides the large number of sensors and subjects, it provides a significant amount of instances of the recorded activities allowing proper assessment of recognition performance. Moreover, subjects were allowed to exhibit large variations in their behaviour making it suitable for the evaluation of activity recognition systems outside of the laboratory.

In addition, we reviewed some of the techniques that have been proposed to cope with changes in the system, including missing data, sensor location/orientation change, as well as the possibility to exploit data from unknown discovered sensors. The Opportunity dataset has been used to test these methods using different network configurations (e.g., using accelerometers only or combining them with gyroscope sensors); while allowing systematic evaluation of the techniques (reported results are average over several repetitions). Moreover, different recognition tasks can also be evaluated; i.e. manipulative gestures in Sections 3.1 and 3.2, and modes of locomotion in Section 3.3. Nevertheless, it has to be mentioned that this type of tests are limited in time since the whole recording for each subject takes place on a single day. Additionally results may vary in scenarios where the user receives immediate feedback about performance, since the feedback may generate changes as the user may try to adapt to the recognition system.

Future efforts should be undertaken to further encourage the use of common datasets and tools. We plan to make this dataset publicly available in the near future and promote the comparison of several recognition approaches from different research groups, probably in the form of an open challenge or competition. Similarly, the use of common tools is also desirable enabling the sharing of data and evaluation methods across the research community. These tools include support for labelling and browsing pre-recorded datasets [23], as well as the software architectures required to support the management of dynamic sensor changes as the ones described here [13].

Acknowledgements

This work is supported by the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, FET-Open grant number: 225938. (For more information: <http://www.opportunity-project.eu>).

References

- [1] J.d.R. Millán, R. Rupp, G.R. Müller-Putz, R. Murray-Smith, C. Giugliemma, M. Tangermann, C. Vidaurre, F. Cincotti, A. Kübler, R. Leeb, C. Neuper, K.R. Müller, and D. Mattia. Combining brain-computer interfaces and assistive technologies: State-of-the-art and challenges. *Frontiers in Neuroscience*, 4:161, 2010.
- [2] M. Bächlin, M. Plotnik, D. Roggen, N. Giladi, J. M. Hausdorff, and G. Tröster. A wearable system to assist walking of parkinson s disease patients. *Methods Inf Med*, 49(1):88–95, 2010.
- [3] M. Tentori and J. Favela. Activity-aware computing for healthcare. *Pervasive Computing, IEEE*, 7(2):51–57, 2008.
- [4] Daniel Roggen, Alberto Calatroni, M. Rossi, T. Holleczeck, Kilian Förster, Gerhard Tröster, Paul Lukowicz, David Bannach, , Gerald Pirkl, Alois Ferscha, Jakob Doppler, Clemens Holzmann, Marc Kurz, Gerald Holl, Ricardo Chavarriaga, Hesam Sagha, Hamidreza Bayati, Marco Creatura, and J. Del R. Millán. Collecting complex activity data sets in highly rich networked sensor environments. In *Seventh International Conference on Networked Sensing Systems*, 2010.
- [5] 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.

- [6] S. Intille, K. Larson, E. Tapia, J. Beaudin, P. Kaushik, J. Nawyn, and R. Rockinson. Using a live-in laboratory for ubiquitous computing research. In *Proc. Int. Conf. on Pervasive Computing*, pages 349–365, 2006.
- [7] T. van Kasteren, A. Noulas, G. Englebienne, and B. Kröse. Accurate activity recognition in a home setting. In *Proceedings of the 10th international conference on Ubiquitous computing*, pages 1–9. ACM Press, 2008.
- [8] Tâm Huynh, Mario Fritz, and Bernt Schiele. Discovery of activity patterns using topic models. In *Proceedings of the 10th international conference on Ubiquitous computing*, pages 10–19. ACM New York, NY, USA, 2008.
- [9] K. Van Laerhoven, H. Gellersen, and Y. Malliaris. Long-term activity monitoring with a wearable sensor node. In *BSN Workshop*, 2006.
- [10] M. Tenorth, J. Bandouch, and M. Beetz. The TUM kitchen data set of everyday manipulation activities for motion tracking and action recognition. In *IEEE Int. Workshop on Tracking Humans for the Evaluation of their Motion in Image Sequences (THEMIS) at ICCV*, 2009.
- [11] Paul Lukowicz, G. Pirkel, D. Bannach, F. Wagner, A. Calatroni, Kilian Förster, T. Holleczeck, M. Rossi, Daniel Roggen, Gerhard Tröster, J. Doppler, C. Holzmann, A. Riener, A. Ferscha, and Ricardo Chavarriaga. Recording a complex, multi modal activity data set for context recognition. In *1st Workshop on Context-Systems Design, Evaluation and Optimisation at ARCS, 2010*, 2010.
- [12] K. Förster, D. Roggen, and G. Tröster. Unsupervised classifier self-calibration through repeated context occurrences: is there robustness against sensor displacement to gain? In *Proceedings of the 13th International Symposium on Wearable Computing*, pages 77–84. IEEE Computer Society, 2009.
- [13] Marc Kurz, Alois Ferscha, Alberto Calatroni, Daniel Roggen, and Gerhard Tröster. Towards a framework for opportunistic activity and context recognition. In *Workshop on Context Awareness and Information Processing in Opportunistic Ubiquitous Systems at 12th ACM International Conference on Ubiquitous Computing*, Copenhagen, Denmark, September 2010.
- [14] Kazuya Murao, Tsutomu Terada, Yoshinari Takegawa, and Shojiro Nishio. A context-aware system that changes sensor combinations considering energy consumption. In *Pervasive '08: Proceedings of the 6th International Conference on Pervasive Computing*, pages 197–212, Berlin, Heidelberg, 2008. Springer-Verlag.
- [15] Hesam Sagha, José del R. Millán, and Ricardo Chavarriaga. A probabilistic approach to handle missing data for multi-sensory activity recognition. In *Workshop on Context Awareness and Information Processing in Opportunistic Ubiquitous Systems at 12th ACM International Conference on Ubiquitous Computing*, Copenhagen, Denmark, September 2010.
- [16] F. Koushanfar, M. Potkonjak, and A. Sangiovanni-Vincentelli. On-line fault detection of sensor measurements. In *IEEE Sensors*, volume 3, pages 69–85, jun 2002.
- [17] L. Paradis and Q. Han. A survey of fault management in wireless sensor networks. *Journal of Network and Systems Management*, 15(2):171–190, 2007.
- [18] C Alippi and M. Roveri. An adaptive cusum-based test for signal change detection. In *ISCAS*. IEEE, 2006.
- [19] Hesam Sagha, José del R. Millán, and Ricardo Chavarriaga. Detecting anomalies to improve classification performance in an opportunistic sensor network. Submitted.
- [20] Kilian Förster, Pascal Brem, Daniel Roggen, and Gerhard Tröster. Evolving discriminative features robust to sensor displacement for activity recognition in body area sensor networks. In *Proc. 5th Int. Conf. on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP 2009)*. IEEE Press, 2009.
- [21] Hamidreza Bayati, José del R. Millán, and Ricardo Chavarriaga. Unsupervised adaptation to sensor displacement in body-worn activity recognition. Submitted.
- [22] Alberto Calatroni, Daniel Roggen, and Gerhard Tröster. A methodology to use unknown new sensors for activity recognition by leveraging sporadic interactions with primitive sensors and behavioral assumptions. In *Workshop on Context Awareness and Information Processing in Opportunistic Ubiquitous Systems at 12th ACM International Conference on Ubiquitous Computing*, Copenhagen, Denmark, September 2010.
- [23] David Bannach, Kai Kunze, Jens Weppner, and Paul Lukowicz. Integrated tool chain for recording and handling large, multimodal context recognition data sets. In *Proceedings of the 12th ACM international conference adjunct papers on Ubiquitous computing*, Ubicomp '10, pages 357–358, New York, NY, USA, 2010. ACM.