Dynamic Adaptation of Opportunistic Sensor Configurations for Continuous and Accurate Activity Recognition

Marc Kurz, Gerold Hölzl, Alois Ferscha
Johannes Kepler University Linz
Institute for Pervasive Computing
Linz, Austria
{*kurz, hoelzl, ferscha*}*@pervasive.jku.at*

***Abstract*—An ever-larger availability of devices that are attached with different sensing capabilities (e.g., smart phones) shifted the challenge in activity and context recognition from the application specific deployment of new sensors to the utilization of already available devices. Therefore, a system that operates in an opportunistic way has to take advantage of the currently available sensing infrastructure in terms of utilizing sensors in form of ensembles that are best suited to execute a specific activity recognition task. Continuous, stable, and accurate activity recognition can be assured if such a system is able to react in real-time to such dynamics in the sensing infrastructure. In detail, this paper tackles the characteristic application cases where sensors spontaneously appear, disappear and reappear in the sensing infrastructure and evaluates the continuousness and stability of the self-adaption methods within the OPPORTUNITY Framework, which is a reference implementation of an opportunistic activity and context recognition system.**

***Keywords*-Opportunistic sensing; activity and context recognition; self-adaptation.**

I. Introduction

Opportunistic activity recognition is characterized by the fact that sensor systems that gather environmental data to infer people's activities are not presumably known at design time of the system [1]. Actually, the activity and context recognition system that operates in an opportunistic way has to take advantage of the currently available devices in order to execute the recognition task as accurate as possible. Another crucial characteristic is that the recognition purpose (i.e., *recognition goal*) is not fixed at design time of the system, either. This goal can be rather defined by a user or an application at runtime of the system [1]. Subsequently, the currently available sensor devices (i.e., the *sensing infrastructure*) have to be queried, and the set of sensors has to be identified that is able to contribute to the recognition goal. According to such a recognition goal, the system configures *ensembles*, which are (sub-)sets of available sensors that are best suited to execute this specific recognition goal [2]. Since the sensor systems that are involved in such ensembles are not presumably known, the system has to dynamically handle changing sensor environments, which also includes different modalities and types of sensors [3]. These characteristics of an opportunistic activity and context recognition systems allow the identification of application cases (see also [4]): (i) *sensor appears*, (ii) *sensor disappears*, (iii) *sensor reappears*, (iv) *sensor delivers reduced-quality data*, and (v) *sensor learns from other sensors*.

This paper tackles the application cases (i), (ii), and (iii), where sensors spontaneously appear, disappear and reappear in the surrounding sensing infrastructure. The remaining two application cases (iv) and (v), where on the one hand the reduced quality data and on the other hand the enhancement of sensor capabilities at runtime are considered have already been subject of discussion and evaluation in recent publications (see [4] and [5]). The experiments and evaluations that are contained in this paper try to answer the question whether the developed concepts like *sensor self-descriptions* (see [1][4][5]), that describe the sensor from a meta-level with respect to the recognition capabilities for specific goals and that enable the dynamic instantiation of activity recognition chains, and the self-organization concepts that enable the dynamic configuration of senor ensembles [5] ensure a continuous, stable and highly accurate activity recognition. Therefore, the paper utilizes a rich dataset that was recently recorded in a kitchen scenario [6]. This allows a high quality evaluation in a repeatable simulated setting with a publicly available dataset. Furthermore, the same setup is also evaluated in an experiment with physical sensor systems to demonstrate the real-time capabilities of the concepts and the OPPORTUNITY Framework, which is a reference implementation of an opportunistic activity and context recognition system together with the sensor self-description and the ensemble configuration concepts.

The rest of the paper is structured as follows. Section II provides an overview of related work. Section III provides a detailed description of the opportunistic activity recognition approach, whereas the main focus lies on the application cases that build the core for further evaluations. Section IV presents an experimental setup based on a rich dataset and evaluates the stability and steadiness of the OPPORTUNITY Framework in terms of dynamically configuring sensor ensembles according to a recognition goal. The last Section V closes with a conclusion and an outlook.

 ISBN: 978-1-61208-219-6

II. Related Work

Activity and context recognition is a research topic that has been tackled from different groups in the last years and decade(s). Mantyjarvi et al. [7] and Bao and Intille [8], present the principles of activity recognition with acceleration sensors mounted on different parts of the body of subjects. The recognition of human activities with a defined set of sensors, with fixed (on-body) position and location is evaluated and tested with different algorithms (see also [9][10][11] as recent examples). The novelty of this paper is the fact that the sensing infrastructure is not presumably known by the system and thus has to react on spontaneous changes to the availability of the surrounding sensors.

Concerning the characteristic application case of such an opportunistic system where sensors that are involved in the recognition process deliver faulty or quality-reduced data (e.g., due to a low battery level), this anomaly can be detected, as described in [1]. This approach is demonstrated in the OPPORTUNITY Framework, which is a working reference implementation of an opportunistic activity and context recognition system, in [5]. There, a quantitative measure called *Trust Indicator* was used to weight the reduced data quality. The second application case that was already demonstrated within the OPPORTUNITY Framework is the crucial task of autonomously enhancing a single sensor's capabilities by observing an active ensemble. This method (also referred to as *transfer learning*) is described in detail in [12] and demonstrated in a running system and tested with respect to real-time requirements in [4]. The remaining characteristic cases that deal with the spontaneous availability of sensors are the core subject of this paper.

Related work that concerns about activity recognition with spontaneously changing sensing environments is - to our best knowledge - very scarce. Approaches exist where sensors are dynamically selected in order to find an accuracy-power trade off [13], or where the activity recognition chain is ported to a sensor platform and distributed to multiple nodes in a wireless sensor network [14].

Chavarriaga et al. [15] present an approach that is derived from an information theoretic concept, where available sensors are dynamically picked with respect to the expected recognition performance of the sensor aggregation. This approach works with the diversity measurements of possible classifier combinations and suffers from the known problem that a ground-truth in the first place is inevitable. Villalonga et al. [16] present a similar use case, where sensor ensembles are configured and the expected recognition accuracy is predicted. These publications are related in a way that not all sensors that would be available are integrated in a sensor ensemble for a specific recognition task but a subset of them, which is optimized in terms of performance. Nevertheless, this paper tackles the problem of handling spontaneous changes in configured sensor ensembles.

III. Opportunistic Activity Recognition

Opportunistic activity and context recognition arises as new working principle since sensor devices have recently become more and more integrated into the daily life. This shifted the effort from the application specific deployment of sensors for specific recognition tasks to the utilization of already available sensors. Therefore, in preceding publications, we have already introduced the concept of sensor abstractions (see [1][3]), which enables the common usage of material as well as immaterial devices as general type *sensor*. This is an important feature in an opportunistic system, since the sensor type and modality cannot be predefined and thus has to be able to handle open-ended sensor environments. The second important method that has already been introduced in recent papers (see [1][4][5]) is the concept of sensor *self-descriptions*. The standardized XML documents are key components in an opportunistic system and provide a two dimensional description of a sensor. The description consists of (i) a technical part that describes the physical working characteristics from the sensing device (could be seen as transcript from the technical specification), and of (ii) a dynamic part that encapsulates the sensor capabilities in terms of recognizing activities. Both parts of the sensor self-description are composed of *SensorML* [17], which is an approved standard. Besides inevitable parts like identifiers, information about the sensor's position, and other relevant information, the dynamic part of the sensor self-description contains key elements that enable the handling of dynamically changing sensor ensembles: *ExperienceItem*s (see also [1][4][5]). The *ExperienceItem*s contain all required building blocks for a dynamic configuration of an activity recognition chain for the dedicated sensor (i.e., (i) *FeatureExtraction*, (ii) *Classifier*, (iii) the accompanying *Classifier Model*, and (iv) the expected/estimated accuracy in form of the *Degree of Fulfillment*). *ExperienceItem*s are highly dynamic elements and can even be added or modified by the OPPORTUNITY Framework at runtime of the system. The application case *"Sensor Learns"*, where a sensor learns the recognition of a specific activity at runtime was demonstrated in [4]. The sensor is able to preserve experience (thus the name *ExperienceItem*), which can be gathered by training the required machine-learning technologies at system's runtime by comparing the sensor readings with the label output of available sensors that are configured in ensembles. In [5], the application case *"Sensor delivers faulty Data"* was demonstrated, which also relies on the capabilities of *ExperienceItem*s to be dynamically updated at runtime. This paper discusses and evaluates the remaining application cases, *"Sensor appears"*, *"Sensor disappears"*, and *"Sensor reappears"*.

Especially the dynamic description with the *ExperienceItem*s enables the dynamic configuration of ensembles according to a recognition goal and permits the required

 ISBN: 978-1-61208-219-6

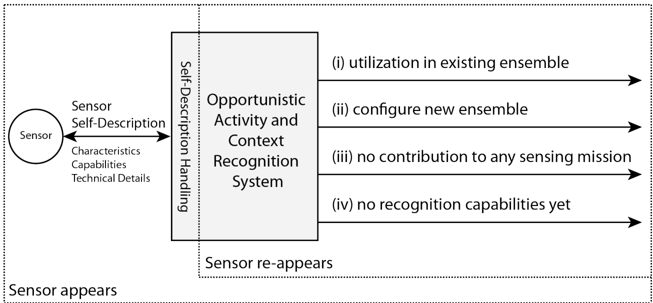

Figure 1. Illustration of the system reaction to the two application cases (i) *sensor appears* and (ii) *sensor disappears*.

stability, continuity and adaptability of the opportunistic activity recognition. Whenever a new sensor appears in the sensing environment, its accompanying (dynamic) self-description is queried and read to be aware about the sensor's present capabilities and its utilization in currently active ensembles. Upon the capabilities according to the dynamic description, the system then decides on the further utilization of the sensor, whereas four different possibilities can be distinguished, as summarized in Figure 1: (i) the sensor is integrated in a running ensemble, (ii) a new ensemble is configured containing the newly appeared sensor, (iii) the sensor's current capabilities do not match an active ensemble, thus the sensor cannot be used, and (iv) the sensor yet does not have any recognition capabilities (i.e., *ExperienceItem*s). Whenever the sensor is not integrated in an ensemble, because the capabilities are not sufficient, or not existing, the sensor could be a candidate to enhance its capabilities by applying *transfer learning* [4][12]. In the case the sensor was already online in the sensing infrastructure, thus is already known by the system on re-appearance, the querying and parsing of the self-description is obsolete. Nevertheless, the four consequent options of further utilization are equal to the aforementioned case (see also Figure 1).

When a sensor disappears - which is the third application case of interest for this paper - three different subsequent system steps have to be distinguished: (i) the sensor was not involved in an active ensemble, therefore no further action is necessary, (ii) the sensor was the exclusive component of an ensemble, thus the execution of the recognition goal cannot be continued, and (iii) the sensor was part of a bigger ensemble, there the execution can be continued with a (probably) reduced recognition rate (dependent on the disappearing sensor's performance). These application cases, which tackle the spontaneous availability of sensors in an opportunistic activity and context recognition system ((i) *sensor appears*, (ii) *sensor disappears*, and (iii) *sensor reappears*) are tested and evaluated in the OPPORTUNITY Framework in an experimental setting that relies on a pre-recorded dataset [6] where physical, on-body mounted sensor devices (see next Section IV) are used.

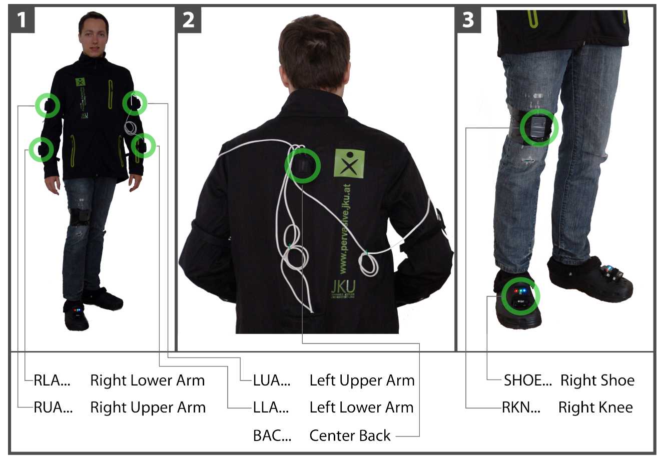

Figure 2. The on-body sensors for the experiment. The setting is similar to [4].

IV. Experiment Setup and Evaluation

The OPPORTUNITY Framework [1] is a reference implementation that realizes the aforementioned concepts of sensor abstractions and sensor self-descriptions. Furthermore, it is capable of executing activity recognition in real-time by applying different sensors, or to process repeatable simulation runs with pre-recorded sensor data. Therefore, the OPPORTUNITY dataset was recorded [6], which combines 72 sensors with 10 different modalities mounted on the body of persons, on objects (e.g., cup, knife, etc.) and in the environment (e.g., fridge, drawer, etc.) in a kitchen scenario. Each of these sensors can be replayed in the OPPORTUNITY Framework as *PlaybackSensor* [3], therefore act as it would be actually available. This allows the generation of repeatable simulation scenarios for evaluation and testing purposes. The experimental setting for this paper is meant to demonstrate the three application cases that tackle the sensors' spontaneous availability (*(re-)appear* and *disappear*). The chosen sensors are all mounted on the body of the test person. The upper body is equipped with 5 sensors of type *Xsens MTx*, which provides 3D acceleration. The sensors are located on both arms (right/left lower/upper arm - *RLA, RUA, LLA, LUA*), and on the upper back (*BAC*). The lower body is equipped with a self-composed bluetooth acceleration sensor on the right knee (*RKN*) and a *SunSPOT* (Small Programmable Object Technology) sensor on the right instep of the foot (*SHOE*). All sensors were operating with a sampling frequency of 100Hz and delivered 3D acceleration data (x-, y-, z-axes). The on-body sensor placement is illustrated in Figure 2.

The challenge is to demonstrate and evaluate the continuity and steadiness of the opportunistic activity and context recognition system by dynamically adapting to different conditions in the sensing infrastructure due to the sensors' spontaneous availability. This is done by dynamic configu-

 ISBN: 978-1-61208-219-6

Table I
OVERVIEW OF ACCURACIES OF THE SINGLE SENSORS FOR THE RECOGNITION OF THE MODES OF LOCOMOTION.

Sensor	RKN	SHOE	LUA	LLA	RUA	RLA	BAC
Accuracy	0.604	0.698	0.858	0.719	0.769	0.709	0.761

ration of different sensor aggregations in form of ensembles over a certain amount of time. We decided to use a rather simple set of activities since not the handling of complex tasks is the major contribution but the accurate and stable activity recognition even if the sensing environment changes. The activities of interest in the exemplary scenario are fixed to the modes of locomotion (i.e., *WALK*, *STAND*, *SIT*, *LIE*). Each of the seven involved on-body sensors is trained to recognize these four activities initially by utilizing the available ground-truth in the dataset (as already mentioned, the training could have also been done by applying the transfer learning [4] approach, but to ensure that all seven sensors have the same conditions, the initial training phase was selected). As features the *mean* and *variance* were used, the classification method was set to be the *NCC* classifier (reasons for these choices are (i) the ease of computation and (ii) the potentially good recognition results [18]). The resulting accuracies for the single sensors were calculated by comparing the predicted activity classes to the actual activity classes after the training phase, and are listed in Table I.

Each sensor's self-description was enhanced with an *ExperienceItem* that contains every piece of information (including the feature extraction method, the classifier method together with the required and pre-trained classifier model in form of a JSON file (see [5] for an example), and the sensor position) to dynamically configure the activity recognition chain at system's runtime. The combination of multiple sensors (respectively multiple recognition chains) in the experiment session is done by applying *MajorityVoting* fusion. This rather simple technology does not need to be trained beforehand, thus can be utilized on the fly. The class where most of the classifiers agree on is selected as fusion result. The prediction of the output accuracy is not a trivial task, in general the diversity measurements between the involved classifiers must be known [15]. This is not yet considered in this paper, but is an open point for future work. In this experiment, each sensor that is capable of contributing and recognizing the activities of interest (i.e., the four modes of locomotion) is integrated in the ensemble, thus is part of the fusion method.

The experiment session lasted for exactly 14 minutes. By mediating a change in the sensing infrastructure (i.e., in a simulated session, sensors of type *PlayBackSensor* [3] can be turned on and off, thus simulating a (re-)appearance and disappearance) different ensembles were configured. In Table II, the thirteen occurring ensemble configurations are listed together with their IDs, the involved sensor devices

Table II
OVERVIEW OF THE ENSEMBLE CONFIGURATIONS TOGETHER WITH THE CALCULATED ACCURACIES.

ID	Active Ensemble	Accuray
1	RKN	0.604
2	RKN + LLA	0.630
3	RKN + LLA + SHOE	0.779
4	RKN + SHOE	0.488
5	RKN + SHOE + BAC	0.853
6	RKN + SHOE + BAC + LLA	0.840
7	RKN + SHOE + BAC + LLA + RLA	0.846
8	RKN + SHOE + BAC + LLA + RLA + RUA	0.817
9	RKN + SHOE + BAC + LLA + RLA + RUA + LUA	0.870
10	RKN + BAC + LLA + RLA + RUA + LUA	0.818
11	BAC + LLA + RLA + RUA + LUA	0.861
12	BAC + LLA + RLA + LUA	0.852
13	BAC + LLA + RLA	0.820

and the accuracy. This accuracy value was calculated by comparing the predicted class as ensemble output with the actual class that can be achieved from the ground-truth. Before the simulated experimental run with the changes in the sensors' availability was done, each of the thirteen ensembles was configured manually and the specific session with the sensors was executed. This was necessary to gather comparable confusion matrices for each of the listed ensembles with the (approximately) same amount of sensor samples and predicted activities. These confusion matrices are illustrated in Figure 3. Each matrix is assigned with the ID from the ensemble (as listed in Table II), and the resulting accuracy. This accuracy can be easily computed out of the confusion matrices, since the main diagonal indicates the correctly predicted classes, the other values the wrong classified activity labels. Each confusion matrix contains on the x- and y-axis the activity class numbers (i.e., 1=*NULL*, 2=*STAND*, 3=*WALK*, 4=*LIE*, 5=*SIT*). During each run (i.e., 14 minutes) the confusion matrices were filled with approximately 500.000 activity classes (this makes approx. 600 classes per second). This means at each second, 600 comparisons from the predicted label to the actual ground-truth label were done to get a significant matrix as base for the calculation of the accuracy.

After the preparation was done (all *ExperienceItem*s generated, the accuracies of the single sensors and of the ensembles calculated), the experimental session was conducted. Figure 4 presents an overview of the actual accuracy of the configured sensor ensembles for the recognition of the modes of locomotion. The x-axis indicates the time in minutes, but also the ID of the active ensemble (e.g., at minute three, the available sensors were mediated to be *RKN*, *LLA*, and *SHOE*). The y-axis contains the accuracy of the active sensor ensemble. As shown, the accuracy and thus the recognition continuity are robust against changes

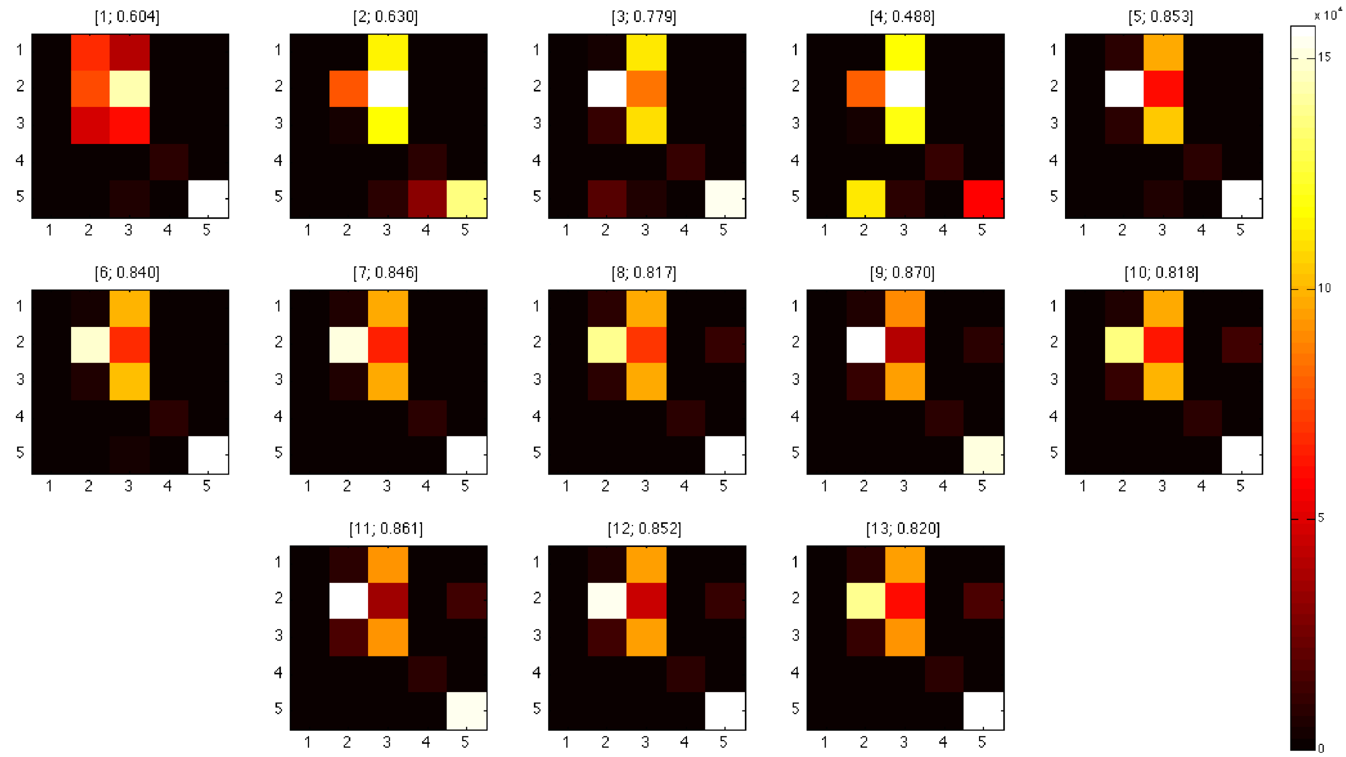

Figure 3. Confusion matrices for the 13 sensor ensembles in the experimental session (the x-axis contains the actual activity class, the y-axis the predicted class). The numbers above each single confusion matrix contain the ensemble ID and the corresponding accuracy compared to the ground-truth.

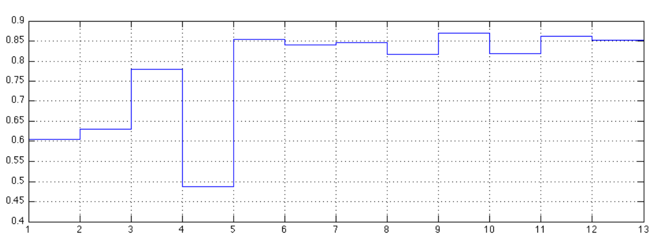

Figure 4. Overview of the accuracies for the occurring ensemble configurations during the experiment session.

in the sensing infrastructure. Based upon the sensor self-descriptions, the system is capable of dynamically adapting to changes in the sensing environment and ensures a continuance of the activity recognition task. This dynamic adaptation to changes in the sensing environment and the resulting stability is a novelty in contrast to conventional activity recognition systems.

V. Conclusion and Future Work

This paper presented the three characteristic application cases *sensor appears*, *sensor reappears*, and *sensor disappears* for an opportunistic activity recognition system. The OPPORTUNITY Framework realizes the concepts of sensor self-descriptions. These XML documents are of highly dynamic nature, since they encode the sensors capabilities in order to recognize certain activities. A sensor can make experience over its life-time (e.g., by manually or autonomous adding of recognition capabilities) and preserve this experience in its self-description in so-called *ExperienceItem*s. Each of these items contains a complete description of an activity recognition chain (i.e., feature extraction, classification and classifier model), together with QoS metrics, like the estimated recognition accuracy. The capability of the OPPORTUNITY Framework to dynamically adapt at runtime to spontaneous changes in the sensing environment is demonstrated in this paper. This enables an accurate, stable and continuous recognition of human activities with dynamically varying sensor settings and can be seen as important building block towards the vision of opportunistic activity and context recognition. This is a novelty in contrast to conventional activity recognition system, since they would fail if sensors disappear, or would not consider new sensors on their appearance.

Concerning future work, one issue is how multiple sensors can be combined to achieve a higher recognition accuracy. The solution so far is to use diversity measures between the sensors. Subsequently, this needs a reliable ground-truth for calculation, which cannot be taken as granted in a real-time application. Nevertheless, to be able to provide an estimation of the accuracy of an ensemble (with the *MajorityVoting* fusion method) these measures are required, since the danger is that multiple sensors classify the same wrong predicted class, which is then taken as winning fusion result. Currently, work is going to calculate these diversity measures at runtime of the system, and preserving this information additionally in the sensor self-descriptions (at least the diversity measures between two sensors without the comparison to the ground-truth).

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] M. Kurz, G. Hölzl, A. Ferscha, A. Calatroni, D. Roggen, G. Tröster, H. Sagha, R. Chavarriaga, J. del R. Millán, D. Bannach, K. Kunze, and P. Lukowicz, "The opportunity framework and data processing ecosystem for opportunistic activity and context recognition," *International Journal of Sensors, Wireless Communications and Control, Special Issue on Autonomic and Opportunistic Communications*, vol. 1, December 2011.

[2] D. Roggen, K. Förster, A. Calatroni, T. Holleczek, Y. Fang, G. Troester, P. Lukowicz, G. Pirkl, D. Bannach, K. Kunze, A. Ferscha, C. Holzmann, A. Riener, R. Chavarriaga, and J. del R. Millán, "Opportunity: Towards opportunistic activity and context recognition systems," in *Proceedings of the 3rd IEEE WoWMoM Workshop on Autonomic and Opportunistic Communications (AOC 2009)*. Kos, Greece: IEEE CS Press, June 2009.

[3] M. Kurz and A. Ferscha, "Sensor abstractions for opportunistic activity and context recognition systems," in *5th European Conference on Smart Sensing and Context (EuroSSC 2010), November 14-16, Passau Germany*, K. K. G. Lukowicz, Paul; Kunze, Ed. Berlin-Heidelberg: Springer LNCS, November 2010, pp. 135–149.

[4] M. Kurz, G. Hölzl, A. Ferscha, A. Calatroni, D. Roggen, and G. Troester, "Real-time transfer and evaluation of activity recognition capabilities in an opportunistic system," in *Third International Conference on Adaptive and Self-Adaptive Systems and Applications (ADAPTIVE2011), September 25-30, Rome, Italy*, September 2011, pp. 73–78.

[5] M. Kurz, G. Hölzl, A. Ferscha, H. Sagha, J. del R. Millán, and R. Chavarriaga, "Dynamic quantification of activity recognition capabilities in opportunistic systems," in *Fourth Conference on Context Awareness for Proactive Systems: CAPS2011, 15-16 May 2011, Budapest, Hungary*, May 2011.

[6] D. Roggen, A. Calatroni, M. Rossi, T. Holleczek, K. Förster, G. Tröster, P. Lukowicz, D. Bannach, G. Pirkl, A. Ferscha, J. Doppler, C. Holzmann, M. Kurz, G. Holl, R. Chavarriaga, M. Creatura, and J. del R. Millán, "Collecting complex activity data sets in highly rich networked sensor environments," in *Proceedings of the Seventh International Conference on Networked Sensing Systems (INSS), Kassel, Germany*. IEEE Computer Society Press, June 2010.

[7] J. Mantyjarvi, J. Himberg, and T. Seppanen, "Recognizing human motion with multiple acceleration sensors," in *Systems, Man, and Cybernetics, 2001 IEEE International Conference on*, vol. 2, 2001, pp. 747 –752 vol.2.

[8] L. Bao and S. Intille, "Activity recognition from user-annotated acceleration data," in *Pervasive Computing*, ser. Lecture Notes in Computer Science, A. Ferscha and F. Mattern, Eds. Springer Berlin / Heidelberg, 2004, vol. 3001, pp. 1–17.

[9] N. Ravi, D. Nikhil, P. Mysore, and M. L. Littman, "Activity recognition from accelerometer data," in *In Proceedings of the Seventeenth Conference on Innovative Applications of Artificial Intelligence(IAAI*, 2005, pp. 1541–1546.

[10] D. Minnen and T. Starner, "Recognizing and discovering human actions from on-body sensor data," in *In Proc. of the IEEE International Conference on Multimedia and Expo*, 2005, pp. 1545–1548.

[11] J. R. Kwapisz, G. M. Weiss, and S. A. Moore, "Activity recognition using cell phone accelerometers," *SIGKDD Explor. Newsl.*, vol. 12, pp. 74–82, March 2011.

[12] A. Calatroni, D. Roggen, and G. Tröster, "Automatic transfer of activity recognition capabilities between body-worn motion sensors: Training newcomers to recognize locomotion," in *Eighth International Conference on Networked Sensing Systems (INSS'11)*, Penghu, Taiwan, Jun. 2011.

[13] 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 *Wireless Sensor Networks*, ser. Lecture Notes in Computer Science, R. Verdone, Ed. Springer Berlin / Heidelberg, 2008, vol. 4913, pp. 17–33.

[14] C. Lombriser, N. B. Bharatula, D. Roggen, and G. Tröster, "On-body activity recognition in a dynamic sensor network," in *Proceedings of the ICST 2nd international conference on Body area networks*, ser. BodyNets '07. ICST, Brussels, Belgium, Belgium: ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2007, pp. 17:1–17:6.

[15] R. Chavarriaga, H. Sagha, and J. del R Millan, "Ensemble creation and reconfiguration for activity recognition: An information theoretic approach," in *Systems, Man, and Cybernetics (SMC), 2011 IEEE International Conference on*, oct. 2011, pp. 2761 –2766.

[16] C. Villalonga, D. Roggen, and G. Tröster, "Shaping sensor node ensembles according to their recognition performance within a planning-based context framework," in *Eighth International Conference on Networked Sensing Systems (INSS'11)*, 2011.

[17] M. Botts and A. Robin, "OpenGIS Sensor Model Language (SensorML) Implementation Specification," OGC, Tech. Rep., Jul. 2007.

[18] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification (2nd Edition)*. Wiley - InterScience, 2001.