

Activity recognition in opportunistic sensor environments

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

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

Alois Ferscha, Marc Kurz, Gerold Hölzl
Institute for Pervasive Computing
Johannes Kepler University Linz
surname@pervasive.jku.at

Hesam Sagha, Hamidreza Bayati, José del R. Millán, Ricardo Chavarriaga
Defitech Foundation Chair in Non-Invasive Brain-Machine Interface
Ecole Polytechnique Fédérale de Lausanne
name.surname@epfl.ch

OPPORTUNITY

OPPORTUNITY is a 3-year long (February 2009-February 2011) EU FP7 project under FET-Open funding¹ with four partners collaborating to develop mobile systems to recognize human activity and context in dynamically varying sensor setups [3]. We envision a mobile system that autonomously discovers available sensors around the user and self-configures to recognize desired activities. It evolves and reconfigures as the environment changes. It self-improves by exploiting re-occurring contexts typical of human behavior and human generated signals related to cognitive states. Overall, it encompasses principles supporting autonomous operation in open-ended environments. OPPORTUNITY thus addresses the limitations of application specific deployed infrastructure. This mainstreams ambient intelligence and improves user acceptance by relaxing constraints on body-worn sensor characteristics, and eases the deployment in real-world open-ended environments. We report on key outcomes of the project during the past two years.

Resilience and adaptation

Real-world deployment of activity recognition systems requires them to be robust against possible changes in the sensor network or user behaviour. This requires the recognition system to be able to self-configure based on the discovery of new sensors or the disappearance or failure of others. In the opportunity project we aim at the development of machine learning techniques that provide these capabilities in conjunction with tools that allow the reconfiguration the sensor network according to the encountered changes. We

¹We acknowledge the support of the commission's research programme under under FET-Open grant number 225938.

have proposed a method for detecting sensors that fail or degrade [5]. It relies on a architecture of classifier fusion and is based on a distance measure between a given sensor and the network overall combined decisions. At operation time anomalous sensors can be automatically removed yielding a graceful performance degradation even when a large proportion of sensors is failing. Alternatively, probabilistic approaches have been developed to handle missing data (e.g. due to transmission problems in wireless networks) [4]. This method infers the missing samples using the conditional distribution between available and missing data sources. This inference can be performed at the raw data level or at the fusion level. In combination these methods can be applied to increase the robustness of real-life deployed systems as they result in highest tolerance to sensor failure without recalibration or external intervention.

Use of unknown new resources

In an opportunistic paradigm, new, unforeseen, resources are likely to be discovered at run-time. This may happen as the environment is upgraded or as the user buys new sensorized gadgets. We devised novel approaches to exploit new resources, even if they are unknown to the system. These methods require no design-time training nor user intervention. This supports the "growth" of a system made of "Context Cells" - sensor nodes capable of activity recognition and local communication. We devised a variation of transfer learning that operates an autonomous transfer of the activity recognition capabilities of one sensor node to another, regardless of the modality and placement of the source and destination nodes. Labels of recognized activities from the first nodes are transferred to the latter, which incrementally associate their sensor signals to the received labels. We devised a novel form of multi-task learning that allows to opportunistically exploit new sensors discovered in the environment to improve the recognition accuracy of an existing system. The approach works by exploiting the natural tendency of relevant activities to form clusters in the feature space. It uses unsupervised clustering of the data of the novel sensors together to disambiguate activities that the initial system may not have been able to distinguish.

Tools and datasets

The development of context aware systems involves a set of complex steps that may differ in detail from case to case but in essence reoccur in most applications. In general the process starts with the collection of a sample data set. Issues that typically need to be addressed include reliable streaming of data into appropriately organized storage, detecting faulty sensors and missing data points, synchronization of sensors and initial data labeling. In most cases the data collection needs to be followed by a post processing enrichment step in which the synchronization points and the labels are refined and missing data points and signals dealt with. This requires different sensor channels to be reviewed, compared with each other and possibly synchronized with a video record of the experiment. For non trivial data sets such post processing can be extremely time consuming. We have found around 10 hours of work to be needed for each hour of recorded data. Once the data has been enriched the core of system development can begin. It is usually an iterative process consisting of feature selection, classifier training and tuning and performance testing. Often the so developed classifier then needs to be ported into a final platform such as a mobile device or a set of such devices.

We have developed an integrated tool chain to support the above process and make it more efficient. The tool chain is built around our CRN Toolbox which is a widely used modular rapid prototyping platform for collecting and processing sensor data. We have implemented (1) a GUI based tool for dynamic monitoring of data flowing through a toolbox application which is crucial to ensure data integrity during long term recordings with a large number of sensors, (2) a data base system based on Apache CouchDB in order to allow sensor data, annotations and accompanying videos to be stored and accessed in an organized way, (3) a GUI based labeling tool for the inspection, annotation, and manual resynchronization of context data stored in the above data base, (4) a trace generation tool that allows combinations of signals from a subset of sensors referring to certain ground truth events to be easily retrieved from the data base and streamed for system training, testing or demonstration.

The OPPORTUNITY Framework

As opportunistic applications draw from the characteristic to use sensing devices according to a recognition goal that just happen to be available instead of defining them at design time, new challenges arise with the implementation of opportunistic activity recognition systems. Therefore, we are currently working on a prototypical implementation of a mobile opportunistic activity recognition system (also referred to as *OPPORTUNITY Framework*), which can be taken as first step towards a ready-to-use middleware for building opportunistic activity recognition applications for different domains. The framework integrates multiple machine-learning technologies (e.g. classification, fusion, feature-extraction, anomaly detection, ...) and by applying self-describing sensors together with a rich domain-knowledge in form of an ontology it is able to choose and configure the set of sensing devices that is best suited to execute a recognition goal (the *ensemble*). Neither the recognition goal, nor the available sensing devices have to be pre-defined, the system reacts at runtime on (i) recognition requests (by translating this request to a machine-readable format - *the goal description*

language), and (ii) on changes in the sensing infrastructure by re-configuration of configured sensing ensembles. The framework is implemented using Java and the OSGi module system and acts as a runtime environment together with a code base and libraries, able to execute autonomously on a target platform. Reasons for implementing the framework in OSGi are (i) the universality of code deployment, (ii) the life cycle management capabilities of OSGi, (iii) the modularity and component oriented paradigm, (iv) the portability and thus compatibility with various different hardware platforms, and (v) the ability to install, restore, start and stop applications at runtime. Generally, OSGi is a collection of bundles that can be interconnected using the *OSGi Wire Admin* service specification, following the producer consumer paradigm to query and propagate data internally from sensors (producers) to sensing mission(s) (consumers) or even from one sensor to another sensor if a direct communication is necessary (for further details be referred to [1] and [2]).

1. CONCLUSION

The highly diverse way in which identical goals unfold in terms of motor actions remains challenging to current activity recognition systems. Ongoing work aims at further method improvement and also seeks an interdisciplinary view on action perception. For instance, understanding the evolved cognitive processes may be key to more robust recognition systems. Welcome community input along these lines.

2. REFERENCES

- [1] M. Kurz and A. Ferscha. Sensor abstractions for opportunistic activity and context recognition systems. In K. K. G. Lukowicz, Paul; Kunze, editor, *5th European Conference on Smart Sensing and Context (EuroSSC 2010)*, November 14-16, Passau Germany, pages 135-149, Berlin-Heidelberg, November 2010. Springer LNCS.
- [2] M. Kurz, A. Ferscha, A. Calatroni, D. Roggen, and G. Tröster. Towards a framework for opportunistic activity and context recognition. In *Proc. of the Opportunistic Ubiquitous Systems Workshop, part of 12th ACM Int. Conf. on Ubiquitous Computing*, 2010.
- [3] D. Roggen, K. Förster, A. Calatroni, T. Hollecsek, Y. Fang, G. Tröster, P. Lukowicz, G. Pirkel, D. Bannach, K. Kunze, A. Ferscha, C. Holzmamm, A. Riener, R. Chavarriaga, and J. del R. Millán. Opportunity: Towards opportunistic activity and context recognition systems. In *Proc. 3rd IEEE WoWMoM Workshop on Autonomic and Opportunistic Communications*, 2009.
- [4] H. Sagha, J. d. R. Millán, and R. 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.
- [5] H. Sagha, J. d. R. Millán, and R. Chavarriaga. Detecting anomalies to improve classification performance in an opportunistic sensor network. In *7th IEEE International Workshop on Sensor Networks and Systems for Pervasive Computing, PerSens 2011*, Seattle, March 2011.