

An Experiment in Hierarchical Recognition of Group Activities using Wearable Sensors

Dawud Gordon¹, Jan-Hendrik Hanne², Martin Berchtold², Takashi Miyaki¹,
and Michael Beigl¹

¹ Karlsruhe Institute of Technology, Karlsruhe, Germany
`firstname.lastname@kit.edu`

² Technische Universität Braunschweig, Brunswick, Germany
`j-h.hanne,m.berchtold@tu-bs.de`

Abstract. Pervasive computing envisions implicit interaction between people and their intelligent environments instead of individual devices, inevitably leading to groups of individuals interacting with the same intelligent environment. These environments must therefore be aware not only of user contexts and activities, but the contexts and activities of groups of users as well. This poster will demonstrate an experiment conducted towards understanding hierarchical multi-user group activity recognition using wearable sensors. The experiment will explore different data abstraction levels in terms of recognition rates, power consumption and wireless communication volumes for the devices involved.

Keywords: group activity recognition, context recognition, distributed systems, multi-user, wearable

1 Introduction

Context and activity recognition provide intelligent devices in the environment with the ability to act proactively in the interest of users. Implicit pro-active interaction based on situational awareness becomes more important in order to prevent us from entering a state of permanent distraction and informational overload due to constantly having to administrate and respond to the myriad of intelligent devices in our immediate environment. As intelligent devices pervade our everyday environment, devices will interact with multiple users at once, meaning that these devices will need to be aware not only of the activities and contexts of individual users, but also of groups of users as well.

The group activity is not necessarily the same as the sum of the activities of the individuals in it [4]. This implies that the activity or context of a group is a function of the activity or context of each individual in the group, in the same way that the context or activity of a single user is a function of sensor measurements. Since carryable or wearable technology is steadily on the rise, leveraging these devices for sensing seems to be a logical approach. Using wearable technology (badges, mobile phones, coffee cups, etc.) for group activity or

context recognition is inherently a hierarchical problem, where data from wearable sensors on multiple users must be aggregated in order to infer the group context. As with other areas of research, such as context prediction [5], group activity recognition (GAR) can be conducted on data at different levels of context abstraction. The lower the abstraction level (e.g. sensor measurements), the more information is contained in the data, but the more energy is required for transmitting this data. Conversely, the higher the data is abstracted towards context labels (e.g. “walking”), the less it costs to communicate, but it contains less information [5].

In this work we present an experiment which will research the accuracy/power-consumption trade-off in GAR based on the context abstraction level using a distributed wearable sensing system. The study is done in an office setting using intelligent coffee cups along the lines of the MediaCup [1] and mobile phones. Using the system, recognition of office group activities such as meetings, presentations and coffee breaks is examined by monitoring individual colleagues.

2 System Design

The system used was made up of a wireless sensor network and a mobile phone. Wireless sensor nodes equipped with 3D acceleration sensors were attached to coffee mugs (Smart Mugs) in a university/office setting. One node was connected to a smart-phone using a USB adapter and acted as a bridge between the WSN and the smart-phone. The nodes sampled activity and context data at the mugs, processed this data to the desired local abstraction level, and forwarded this data to the smart-phone for further processing to classify the group activity.

The wireless sensor nodes used in this experiment were jenParts from the open-source Jennisense Project ³ which uses a combination of the ConTiki operating system [3] and the JENNIC JN5139 wireless microprocessor and includes open source sensor and bridge board designs. The nodes were equipped with a battery and an analog 3D acceleration sensor and were attached to coffee mugs, which we will refer to as Smart Mugs (see Fig. 1).

The nodes sample the sensors at a rate of 33 Hz and can be configured to process this data to different context abstraction levels. The preprocessing operations are dependent on the mode: in raw sensor data mode, the sensors are sampled but no further preprocessing operations are performed (low context abstraction level [5]). In feature mode, the features average and variance are calculated over the window (medium abstraction). In classification mode, these features are classified to local activities using a previously trained local classifier (either kNN or DT, high abstraction), the result of which indicates which activity is being performed by the user holding the mug. This locally generated information is finally forwarded to the smart-phone which tries to identify the global group activity based on this local information.

³ The Jennisense project: <http://www.jennisense.teco.edu>

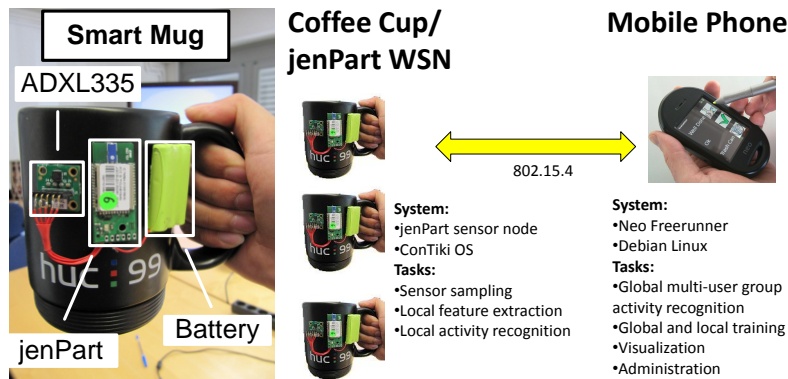


Fig. 1. Left: The Smart Mug **Right:** Network Topology of the Experiment

The smart-phone used is the Neo Freerunner (Neo) from the Openmoko Project⁴ running a distribution of Debian Linux. It serves as a processing platform for classifying the global context or activity based on the data aggregated from all nodes in the WSN, either raw sensor data, locally extracted features, or locally classified activities depending on the mode. The Neo also serves as a context classifier training platform for the jenParts in the WSN. Following the approach presented by Berchtold et al. [2], after being set in training mode by the Neo, each jenPart gathers data and forwards it to the Neo along with a local annotation indicated by labeling activities using the button on the jenParts. Once this process is complete, the Neo trains the selected classifier, segments the trained classifier into packet-sized chunks, and sends these chunks sequentially to the nodes in a JSON format.

3 Experiment

In order to research the energy consumption and quality of recognition based on the context abstraction level in a distributed scenario, a multi-user experiment was conducted. In the first phase, activities were performed to train the local classifiers on the Smart Mugs. In the second phase, the three subjects conducted other activities as a group in order to train and test the global group classifier. The power consumption of each node as well as the consumption of the mobile phone was also monitored. During the course of this experiment, 3 subjects performed 7 different activities, 3 of which were group activities and 4 of which were individual activities involving the Smart Mugs. In total, over 45 minutes of data were collected, consisting of over 22,700 sample windows, approximately evenly distributed between group and individual activities.

In the first phase of the experiment, the Smart Mugs were evaluated in terms of power consumption and accuracy for different levels of abstraction. Each user

⁴ <http://www.openmoko.com/>

performed a set of activities, each one for a duration from approx 2 - 15 minutes. The local activities were as follows: the subject has placed the mug on the **table** (or other surface), the subject is holding the mug in their **hand**, the subject is **drinking** from the mug, and the subject is **gesticulating**.

In the second phase the global classification rates for all modes were evaluated. The subjects conducted the following activities for 4 - 5 minutes a piece using the Smart Mugs trained in the previous phase: **Meeting, Presentation** and **Coffee break**. Offline, the first half of the data was used to train 6 different classifiers and the second half to evaluate their performance. For each of the 3 abstraction levels which are transmitted to the Neo, both a kNN and a DT classifier were trained for global GAR for comparison, and the second half of the data was used to evaluate the performance of the 6 different classifiers.

4 Conclusion

This paper described an experiment in the field of multi-user group activity recognition using wearable sensors. A hierarchical approach to recognizing these activities was examined, where nodes forward information about the subject's activities to a central device which attempts to recognize the activity that the group is currently performing. In this experiment, different context abstraction levels were conducted at the mugs and the entire network was monitored in terms of power consumption and activity recognition rates. The results have not yet been evaluated, but they are expected to provide insight into the effects of the context abstraction level on group activity recognition using wearable sensors.

The authors would like to acknowledge funding by the European Commission for the ICT project "CHOSeN" (Project number 224327, FP7-ICT-2007-2).

References

1. Michael Beigl, Hans-W. Gellersen, and Albrecht Schmidt. Mediacups: experience with design and use of computer-augmented everyday artefacts. *Comput. Netw.*, 35:401–409, March 2001.
2. Martin Berchtold, Matthias Budde, Dawud Gordon, Hedda Schmidtke, and Michael Beigl. ActiServ: Activity recognition service for mobile phones. In *ISWC'10: Proceedings of the Fourteenth International Symposium on Wearable Computers*, pages 83–90, Seoul, S. Korea, 2010. IEEE Computer Society.
3. Adam Dunkels, Bjrn Grnvall, and Thiemo Voigt. Contiki - a lightweight and flexible operating system for tiny networked sensors. In *Proceedings of the First IEEE Workshop on Embedded Networked Sensors*, Tampa, Florida, USA, November 2004.
4. Tao Gu, Zhanqing Wu, Liang Wang, Xianping Tao, and Jian Lu. Mining emerging patterns for recognizing activities of multiple users in pervasive computing. In *Mobile and Ubiquitous Systems: Networking Services, MobiQuitous, 2009. MobiQuitous '09. 6th Annual International*, pages 1–10, July 2009.
5. Stephan Sigg, Dawud Gordon, Georg von Zengen, Michael Beigl, Sandra Haseloff, and Klaus David. Investigation of context prediction accuracy for different context abstraction levels. *IEEE Transactions on Mobile Computing*, 2011. (to appear).