

University of East London Institutional Repository: <http://roar.uel.ac.uk>

This paper is made available online in accordance with publisher policies. Please scroll down to view the document itself. Please refer to the repository record for this item and our policy information available from the repository home page for further information.

To see the final version of this paper please visit the publisher's website. Access to the published version may require a subscription.

Author(s): Naeem, Usman; Bigham, John.

Article title: Recognising Activities of Daily Life through the Usage of Everyday Objects around the Home

Year of publication: 2009

Citation: Naeem, U. and Bigham, J. (2009) "Recognising Activities of Daily Life through the Usage of Everyday Objects around the Home" in Proceedings of the 3rd International Conference on Pervasive Computing Technologies for Healthcare, Technologies to Counter Cognitive Decline Workshop, London, March 2009 pp.1 - 4

Link to published version:

<http://dx.doi.org/10.4108/ICST.PERVASIVEHEALTH2009.6059>

DOI: 10.4108/ICST.PERVASIVEHEALTH2009.6059

Recognising Activities of Daily Life through the Usage of Everyday Objects around the Home

Usman Naeem

School of Electronic Engineering and Computer Science
Queen Mary University of London
London, United Kingdom
usman.naeem@elec.qmul.ac.uk

John Bigham

School of Electronic Engineering and Computer Science
Queen Mary University of London
London, United Kingdom
john.bigham@elec.qmul.ac.uk

Abstract—The integration of RFID sensors into everyday products has become a widespread solution for increasing efficiency in supply chain management. This has also led to a way of being able to monitor everyday activities in the home based on when and how these products are used, which is less intrusive than other monitoring approaches such as visual based systems. Monitoring activities in a home environment can be seen as a good way of analyzing behavior and tracking functional decline among elderly people. This paper describes a hierarchal approach for activity recognition using object usage data generated by everyday products used around the home. The motivation of this work is to allow people with early Alzheimer’s disease to have additional years of independent living before the disease reaches a stage where the person is fully dependable on someone else.

Keywords- Hierarchal Activities of Daily Life; Alzheimer’s Disease; Task Sequences; Object Usage;

I. INTRODUCTION

Life expectancy for people in Europe has steadily been increasing, which has led to more elderly people in the society. It is also predicted that 25% of the European population will be made up of people aged over 65 [1]. This also leads to a concern, as the health of elderly people also tends to deteriorate as their age increases. Alzheimer’s disease is a common impairment found in the elderly population, which currently cost the UK alone and estimated £17 billion a year. With the elderly population on the increase it also predicted that by 2025 there will be over one million people in the UK alone who will be suffering from dementia [2]. Leading an independent life can be very difficult for people who are in the latter stages of this disease, as they find it difficult to make decisions themselves and are dependent on the person who is taking care of them. In addition the structure and demands of society make it difficult for children to look after or provide assistance to their aging parents who suffer from this disease. This is normally due to lifestyle preferences and commitments, as well as geographical mobility with children working and living remotely from their parents.

A form of assistance that is currently given to patients with Alzheimer’s disease is the regular visits from carers and health visitors [3], who prescribe a set of Activities of Daily Life (ADL) in order to deal with forgetfulness as well as providing

the elderly person with stimulation and a framework for an independent life [4]. However there can sometimes be a situation where the elderly person has regular lapses in memory and forgets what the activity that they were supposedly doing. This then leads to frustration and anxiety for the elderly person who becomes aware that they are slowly losing their independence. Being able to recognise ADLs not only provides information of the elderly person being safe, it also enables the possibility of being able to provide assistance given a particular situation, e.g. if an elderly person has forgot what activity they were conducting. In addition the recognition of activities can provide useful information about the ADL and what they are meant to be doing next, or even provide alternative options.

When conducting activity recognition it is important to collect features regarding the activity in an unobtrusive manner, which does not invade the privacy of the person being monitored, namely the use of visual systems. Two of the most common approaches that detect features without being intrusive are the use of Radio Frequency Identification (RFID) and wearable sensors. ‘Dense sensing’ [5] has become a favoured technique for detecting features with RFID, which is based on tagging numerous objects around the home (e.g. Kettle) with wireless transponders and sensors that transmit information whenever an object is used or touched via an RFID reader. This technique is popular due to the transponders being durable, re-usable, small size, low cost and easy to install. However, deploying a large number of transponders can sometimes be tedious to install, as well as that the transponders do not function properly with objects made of metal. The use of wearable sensors are also popular for feature detection, which have been used in the form of accelerometers [6] and audio sensors that provide data about particular body motions [7] and the surroundings where the activity is being conducted. Wang et al [8] have shown that a range of fine grained arm actions like ‘chopping with a knife’ may be determined by using feature detection technique based on wearable sensors. The identification of these actions (e.g. arm movement) can then be combined [9] with object data (e.g. clothes iron) from RFID transponders in order to achieve accurate ADL recognition. One of the drawbacks of using wearable sensors is the inconvenience that is caused to the person being monitored, as they are required to wear a range of sensors around their body while carrying out every day activities.

Markov models have been a popular choice for the construction of probabilistic models for carrying out activity recognition from the features detected, one such approach was by Wilson et al [10] where task recognition experiments were conducted and analysed by Hidden Markov Models (HMM) based around Viterbi algorithm. This was to determine which task is currently active from a sequence of objects used to perform an activity. This and similar approaches are not as efficient when the tasks can be carried out in a random order. This is a problem as human beings often vary the order of task execution when achieving a goal. The use of ontologies [11] and data mining [12] techniques have also been applied in order to solve the problem of missing data and incomplete feature data (e.g. missing objects from a sequence). The ontologies are used to build reliable activity models that are able to match between unknown objects with a word in an ontology which is related to the object in the activity models. For example, a Cup object could be substituted by a Mug object in the activity model ‘Make Tea’ as it uses Mug.

The work in this paper is based around activity recognition through object usage data collected using RFID sensors. Extensive monitoring can sometimes be seen as intrusive and affect people’s privacy therefore our approach utilises more knowledge about the structure of the ADLs as opposed to simply relying on a large number of objects needing to be tagged or labelled. The automation element of this approach is based on plans (representing ADLs) structured hierarchically where knowledge at different levels is used to recognise the activity. In addition the approach is able to analyse the intentions of the elderly person, by being able to predict what ADL they may carry out next. This type of analysis allows the platform to provide assistance while an elderly person with Alzheimer’s conducts an ADL, as well as instituting safeguards.

II. MODELLING ACTIVITIES OF DAILY LIFE

For the work in this paper, ADLs have been modelled in a hierarchal structure, which allows knowledge at different levels of abstraction to be decomposed into subcomponents for reliable activity recognition. Within the proposed hierarchal structure, ADLs are modelled as plans. These plans can contain sub-plans, which can be nested within one or more ADLs. When a plan cannot be decomposed any further it is then known as a task. When performed, a task generates sensor events based on the objects used to perform the activity, e.g. a kettle being used to make tea, and so task recognition is based on analysing sensor data, while ADL recognition is based on recognising the constituent tasks. The ADLs are represented in a hierarchal plan representation language called Asbru. This is a task-specific and intention-oriented plan representation language that was developed as a part of the Asgaard project to represent clinical guidelines and protocols in XML [13]. Asbru has the capability to allow each skeletal plan to be flexible and to work with multiple skeletal plans, this is useful as unknowingly people can sometimes carryout multiple ADLs.

Figure 1 shows an example of a Hierarchal ADL (HADL) for the activity ‘Make Breakfast’, which contains a simple sequence of tasks: ‘Make Tea’ and ‘Make Toast’. The sequences of the sensor events at the lowest level (‘Kettle

Sensor’, ‘Fridge Sensor’, ‘Tea Bag Bowl Sensor’, and ‘Sugar Bowl Sensor’) correspond to sensors triggered during the task ‘Make Tea’ but these may be in any order, or be performed in parallel. Between the sensor events and tasks tier is the Task Associated Sensor Events (TASE) tier, where the sensor events are mapped to their associated hypothesised tasks. Each sensor event associated with the task that makes use of the object is mapped as an TASE. For example in Figure 1, ‘Kettle’ sensor event can be associated with ‘Make Tea’ or ‘Make Coffee’. If the tasks are denoted by letters so

- Task “Make Tea” is denoted by letter A
- Task “Make Coffee” is denoted by letter B

Then the sensor event “Kettle Sensor’ is replaced by Make Tea| Make Coffee = A+B, where + is used to represent the disjunction.

Once the sensor events have been mapped into the associated tasks they are then partitioned into segments where each segment is mapped to a task, from which the activity recognition and intention analysis is carried out.

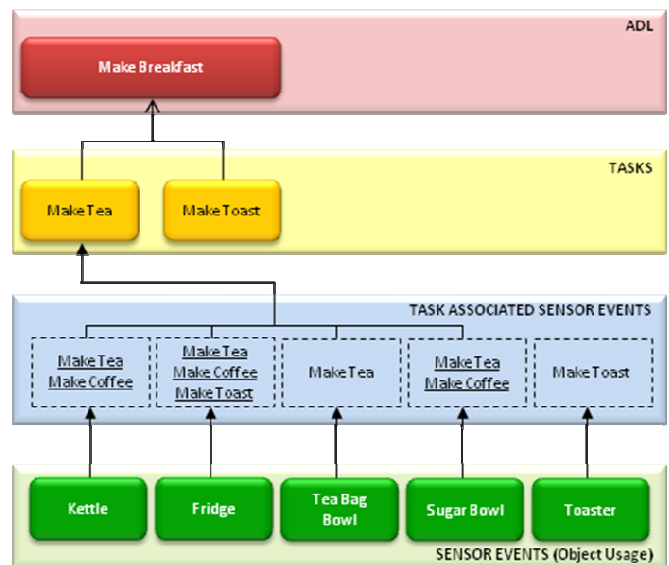


Figure 1. An hierarchal representation of ADL ‘Make Breakfast’

III. TASK RECOGNITION

Recognising tasks can be simple process of segmenting sensor events into segments that correspond to a particular task. However one of the deficiencies with this approach is that there is always a possibility that the sensor event segments that have been generated might be incorrect as they do not bear any resemblance with the task that is actually being carried out. Our task recognition approach assigns a probability $P[b | a]$, where a is a task and b a sensor event. These are established during a training data phase or assigned as prior probabilities. Using the recognition from the higher tier of our approach it is possible determine the probability proportions of $P[a | b]$ given the activity that has been recognised in the higher tier.

We have developed three different approaches to task recognition. One is based on Multiple Behavioural Hidden Markov Models (MBHMM) [14] and the other using a technique inspired from an approach for text segmentation [15]. The third approach is based on Generating Alternative Task Sequences (GATS) from a stream of object usage data based on the product of each task associated sensor event. The work in this paper describes how the combination of the GATS approach and the higher tier plans are able to achieve reliable ADL recognition and analysing the intentions of the person conducting the activity.

The GATS approach is used to provide ordered lists of alternative task sequences given an input set of sensor events. Each of these task sequences has an associated cost, where the cheapest task sequence is taken as the most likely task sequence as the cost function is intended to reflect the compliance of the task sequence with the event sequence and the relative frequencies of ADLs in the higher tier.

The function of the GATS approach can be represented as:

$$e_1, e_2 \dots e_n \rightarrow \{ \langle TS_1, c_1 \rangle + \langle TS_2, c_2 \rangle + \dots + \langle TS_m, c_m \rangle \} \quad (1)$$

where e_n represents the sensor events in the order of observation, TS stands for a task sequence consistent with the event sequence. m is a parameter chosen when the task recogniser is asked for its set of task sequences that match the events. A reason for doing this is to limit the number of possibilities generated, as m is treated as an upper limit, meaning if there are fewer than m possibilities then only the actual possibilities are generated. For example, after the events e_1 , e_2 and e_3 are observed, a list of two possible task sequences, ABC and ABD might be generated, where A , B , C and D are tasks. ABC will have a cost and so will ABD . The set of alternative and mutually exclusive task sequences as well as their costs will be represented as $\{ \langle ABC, c_1 \rangle + \langle ABD, c_2 \rangle \}$.

If it is not evident what task is being conducted from the current task sequences, the higher tier can request from the GATS approach to provide further n task sequences, it will generate an additional $\{ \langle TS_{m+1}, c_{m+1} \rangle + \dots + \langle TS_{m+n}, c_{m+n} \rangle \}$ task sequence.

As a new sensor event is detected the task recogniser is invoked, which then computes a new set of task sequences. Making this the output of this new invocation relate to the previous is a function of the GATS approach, as it has the ability to recognise that the more recent tasks are more important when computing cost function associated with a task sequence. This is made possible by some exponential weighting of costs, where the matching of tasks to the more recent sensor events is given more weight.

As the GATS approach takes into consideration all the possible types of task sequences given the task associated sensor events, it therefore mitigates the chances of not being able to recognise tasks that have been carried out via different variations.

IV. ADL RECOGNITION

The high level ADL recognition gives an overview of all the possible ADLs that could occur within a given time frame. In addition, the higher tier is capable of taking into consideration any overlapping ADLs as well as being able to differentiate which ADL is currently active from the tasks that are discovered in the lower level task recognition. The higher tier recognition component takes a task sequence for input, and creates as output a list of alternative ADLs sets, each with an associated utility.

An ADL set is a collection of probable ADLs that are generated given the tasks sequence that have been recognised and generated by the lower tier GATS approach. The utility of each these ADL sets is based on the cost of each task sequence. The term sequence is not used for ADLs as some of the ADLs can be concurrent. Events and tasks, however, are considered atomic and so the term event sequence and task sequence is valid. So each task sequence $t_1, t_2 \dots t_m$ generates the alternatives $ADLS_1, ADLS_2 \dots$ where $ADLS_i$ denotes a set of ADLs consistent with the task sequence. 1 is the utility of the ADL set.

$$t_1, t_2 \dots t_m \rightarrow \{ \langle ADLS_1, \rho_1 \rangle + \langle ADLS_2, \rho_2 \rangle + \dots + \langle ADLS_m, \rho_m \rangle \} \quad (2)$$

The utility of each ADL set is based on the cost of each segmented task sequence. Therefore in order to achieve accurate ADL recognition it is significant to recognise as many tasks as possible within a window of events. However, it can be difficult to generate the utility of every possible ADL in the library at the time of an activity being conducted. Consequently, the utilities generated for each ADL set are based on ADL schedules within a certain time frame (example 9.00am to 9.15am). This makes the recognition process more manageable and accurate by eliminating some of the unlikely possibilities at the very outset of the recognition process. These ADL schedules are inspired by real life planned activity examples constructed by the Alzheimer's Association for people with dementia. The mission of the association is to help people who suffer from dementia with an organised day consisting of activities that are modelled to meet each individual's preference, as well as that the objective of these planned activities is to enhance the individual's self esteem and improve quality of life by providing them with purpose and meaning to their life [3]. These activities are split into different time segment throughout the course of the day, hence our construction of the ADL schedules is based on this interval-based structure. However there are also certain ADLs that can occur at any time (e.g. phone ringing), which we refer to as interruption ADLs and these are modelled within each ADL schedule in the ADL library. This is all made possible by the representation language 'Asbru', as it can represent and model timing intervals between ADLs.

Asbru is a plan representation language based in XML, each ADL is constructed in XML. Therefore when constructing an ADL it is possible to construct one ADL per XML file, or it is even possible to construct a series of ADLs into one larger XML file, (e.g. ADL schedules). Both of these ways can lead to a situation where the XML file will contain the same tasks

that belong to different ADLs. For example, task ‘Make Tea’ could belong to ADL ‘Make Breakfast’ and ‘Daytime Snack’. If there is an instance where there are two possibilities then this is represented by two paths specifying the location of the task that has been detected. In order to distinguish between the different possibilities and correctly determine which activity is currently being conducted is done by calculating discrepancies and surprise indexes.

A discrepancy is a task that has not been detected, which should have been detected when the ADL was executed. The overall discrepancy of an ADL is computed by summing the discrepancies of the sub-activities. This is further aided by a surprise index, which is used to account for the fact that the absence of some tasks can be more unusual than others, and enumerates this by accruing a measure of how likely a task is when an ADL is being executed. When an ADL is being performed, if the surprise index exceeds an ADL’s given surprise threshold then it means that the ADL has not been detected correctly or another ADL is being conducted. The surprise index is the maximum of the conditional probability $P[a | b]$ of a missing sub-activity and tasks (a) given the ADL (b) that is being conducted, while a discrepancy is computed whenever there is any missing mandatory task, such as ‘Make Toast’ for the ADL ‘Make Breakfast’.

V. EXPERIMENTS & RESULTS

The objective of this experiment conducted was to see how well this approach of combining the GATS approach with the higher tier of our hierarchal approach is for recognising ADLs and its constituent tasks. This experiment was also conducted to validate the GATS approach by seeing if the relevant tasks were being segmented correctly. The experiment was conducted with non-intrusive RFID transponders installed around a kitchen and on its cupboards and objects, such as kettle, dishwasher, utensils, and toaster. The experiments were based around 5 ADLs, which were made up of a series of sub-activities and tasks.

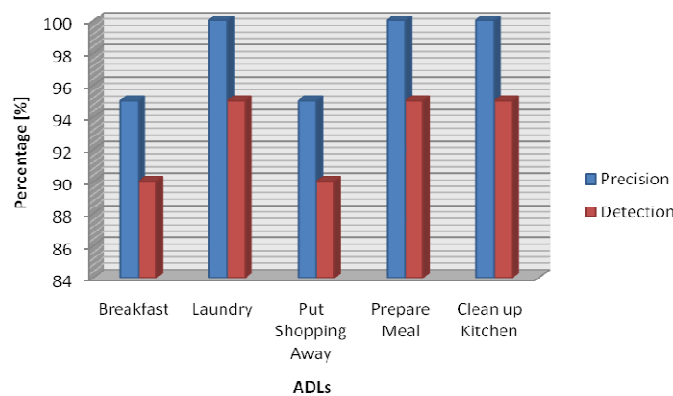


Figure 2. Experiment Results

The results in Figure 2 show that the precision rates were high for all of the ADLs, as this proves that the GATS approach considers all the possible task sequences when carrying out task recognition. However it only takes into

consideration the stream with the lowest cost. Consequently, the stream with the lowest cost provided the segmented tasks which more than often consisted of the relevant tasks that had been conducted. In addition the detection results were high, as the expected tasks and ADLs frequently matched the ground truth data collected.

VI. FUTURE WORK

Currently we are investigating ways to generalise the activity recognition capability of this hierarchical approach outside the framework of the core ADLs constructed to support recognition.

REFERENCES

- [1] W.C. Mann, “The Aging Population and Its Needs,” in IEEE Pervasive Computing, vol.03, no.2, pp.12-14, April-June, 2004
- [2] M. Knapp, “Dementia UK,” London School of Economics and Institute of Psychiatry at King’s College London, Published by Alzheimer’s Society, UK, February 2007
- [3] M.Philipose, “Lare-Scale Human Activity Recognition Using Ultra-Dense Sensing,” in The Bridge, vol.35, no.4, National Academy of Engineering, 2005
- [4] The Alzheimer’s Association, “Activities at Home, Planning the day for a person with dementia,” Alzheimer’s Association National Office, Chicago, pp.1-12, 2005
- [5] M.Philipose, K.P. Fishkin, M. Perkowitz, D.J. Patterson, H. Kautz, D. Hahnel, “Inferring Activities from Interactions with Objects,” in IEEE Pervasive Computing Magazine, vol.3, no.4, pp.50-57, 2004
- [6] L. Bao, S.S. Intille, “Activity Recognition from User-Annotated Acceleration Data,” in Proceedings of the 2nd International Conference on Pervasive Computing, pp.1-17, April, 2004
- [7] M. Ermes, J. Pärkkä, J. Mäntyjärvi, I. Korhonen, “Detection of Daily Activities and Sports With Wearable Sensors in Controlled and Uncontrolled Conditions,” in IEEE Transactions on Information Technology in Biomedicine, vol.12, no.1, January, 2008
- [8] S. Wang, W. Pentney, A. Popescu, T. Choudhury, M. Philipose, “Commonsense-based Joint Training of Human Activity Recognizers,” in Proceedings of the 21st International Joint Conference on Artificial Intelligence, 2007
- [9] M. Stikic, T. Huynh, K. Van Laerhoven, B. Schiele. “ADL Recognition Based on the Combination of RFID and Accelerometer Sensing,” in 2nd International Conference on Pervasive Computing Technologies for Healthcare, January, 2008
- [10] D.H. Wilson, D. Wyaat, M. Philipose, “Using Context History for Data Collection in the Home,” in Pervasive 2005. LNCS, vol. 3468, 2005
- [11] E. Munguia-Tapia, T. Choudhury, M. Philipose, “Building Reliable Activity Models using Hierarchical Shrinkage and Mined Ontology,” in Proceedings of the 4th International Conference on Pervasive Computing, pp.17–32, 2006
- [12] D. Wyatt, M. Philipose, T. Choudhury, “Unsupervised Activity Recognition Using Automatically Mined Common Sense,” in Proceedings of the 20th National Conference on Artificial Intelligence, pp.21-27, July, 2005
- [13] Fuchsberger, C., Hunter, J., McCue, P, “Testing Asbru Guidelines and Protocols for Neonatal Intensive Care” in Proceedings of the Tenth European Conference on Artificial Intelligence, pp.101–110, 2005
- [14] U. Naeem, J.Bigham, “A Comparison of Two Hidden Markov Approaches to Task Identification in the Home Environment” in Proceedings of the 2nd International Conference on Pervasive Computing and Applications, pp. 383-388, 2007
- [15] U. Naeem, J. Bigham, J. Wang, “Recognising Activities of Daily Life using Hierarchical Plans” in Proceedings of the 2nd European Conference on Smart Sensing and Context, LNCS 4793, pp. 175-189, 2007