

A Framework for Semi-Supervised Adaptive Learning for Activity Recognition in Healthcare Applications

Prankit Gupta and Praminda Caleb-Solly

Bristol Robotics Laboratory, University of the West of England, Bristol BS16 1QY, UK
prankit.gupta@brl.ac.uk, praminda.caleb-solly@uwe.ac.uk

Abstract. With the growing popularity of the Internet of Things and connected home products, potential healthcare applications in a smart-home context for assisted living are becoming increasingly apparent. However, challenges in performing real-time human activity recognition (HAR) from unlabelled data and adapting to changing user health remain a major barrier to the practicality of such applications. This paper aims to address these issues by proposing a semi-supervised adaptive HAR system which combines offline and online recognition techniques to provide intelligent real-time support for frequently repeated user activities. The viability of this approach is evaluated by pilot testing it on data from the Aruba CASAS dataset, and additional pilot data collected in the Bristol Robotics Lab’s Assisted Living Studio. The results show that 71% of activity instances were discovered, with an F1-score of 0.93 for the repeating “Meal_Prep” activities. Furthermore, real-time recognition on the collected pilot data occurred near the beginning of the activity 64% of the time and at the halfway point in the activity 96% of the time.

Keywords: Activity Recognition, Smart-home, Healthcare.

1 Introduction

The term smart-home (SH) refers to the concept of integrating everyday items/home appliances with various sensors, actuators, and relays, and their inter-networking to achieve various forms of automation [1]. An increasing population of older adults with ageing-related impairments and long-term conditions is putting a greater amount of pressure on healthcare providers, who are struggling to balance high-quality care with reduced budgets. The utilization of IoT technology within a SM context to provide Ambient Assisted Living (AAL) is regarded as a key solution to achieve this balance [2].

An important area of research in SHs for AAL is human activity recognition (HAR). Offline and online HAR are two different approaches for detecting activities, with online methods working in real-time as the activity is occurring [3]. Most conventional HAR approaches that have shown promising results in the past have used supervised learning, requiring large amounts of expert-labelled user activity training data which is both difficult to obtain and not practical for real-world deployment [4]. These systems

also largely remain unable to adapt to changing user health and behaviour over time, as they require supervised re-training with new activity data to account for this.

This paper proposes an alternative approach to conventional HAR, where unlike traditional systems the focus is not on training a classifier to identify every instance of every activity, but rather on discovering frequently repeating activity patterns. Accordingly, this paper presents a framework which utilizes a pre-defined short-term system memory for performing offline HAR on unlabelled data and updating the user activity model using Bayesian Networks for real-time detection. This combination of offline and online HAR works simultaneously to constantly learn activity patterns and evolve the user activity model on a short-term basis to account for any new or changing activities. This evolving user activity model could be used to set-up automated prompts for users with cognitive decline as well as track and continually assess user health to complement carer support. An implementation of this approach is also pilot tested with the Aruba CASAS dataset [5] and pilot data collected in the Assisted Living Studio (ALSt) which is set up as a realistic home environment in the Bristol Robotics Laboratory.

The rest of the paper is divided as follows – Section 2 provides background information and outlines related work; Section 3 formulates the problem; Section 4 presents the overall framework for the HAR system, including descriptions of the offline and online HAR systems and machine learning methods used; Section 5 describes the data sets used for pilot testing the approach; Section 6 outlines the procedure followed in developing and testing the system; Section 7 presents the results and related discussion; and finally Section 8 summarises the conclusions and discusses future work.

2 Background and Related Work

This section reviews existing research to establish the current state of the art in SH sensor systems, as well as determine the strengths and limitations of the machine learning algorithms and adaptive environments used.

2.1 Smart-Home Sensor Systems

SHs can consist of a variety of different sensors ranging from ambient sensors, body-sensor networks, to video-based solutions [6]. Ambient wireless sensor systems consist of sensors embedded in the environment of the user such as Passive Infrared (PIR) motion, magnetic contact, temperature, and LUX sensors [7]. Generally, a large number of ambient sensors are required to be present in the room in order to track an activity [8], however they are viewed as less intrusive and more acceptable than other types of sensing techniques which include video monitoring [9]. Body-sensor systems include devices that are physically worn or carried by the user such as fitness monitors and smartphones [10]. These can provide physiological data such as heart-rate and body temperature, along with activity data through embedded accelerometer, gyroscope, etc. The downside of body sensors is that they are often viewed as intrusive and older adults may forget to wear them every day [6]. Video-based methods generally provide the

most contextual information for HAR, but are viewed as less acceptable due to security and privacy issues [11].

2.2 Machine Learning Algorithms for Activity Recognition Systems

Offline Learning HAR systems.

In activity recognition, semi-supervised and unsupervised algorithms are generally deployed for discovering activities from unlabelled data in an offline system. Kwon et al. [12] present a comparison of unsupervised clustering techniques which can be used to discover activity clusters in unlabelled sensor data. Their research focused on recognising activities from smartphone data such as walking, sitting, standing and running. The unsupervised machine learning techniques included Gaussian mixture models, Hierarchical Agglomerative Clustering (HAC), DBSCAN, and K-means clustering. HAC and DBSCAN achieved the highest accuracy and the authors noted HAC for its flexibility and ability to discover hierarchies in the data. HAC works by initially assigning each data point as a cluster, then merging the clusters together at each step. A cut-off can be set as to when the merging should stop, and the various steps of the clustering process can be viewed in the form of a dendrogram.

Real-time Learning HAR and Adaptive Environments.

Yala et al. [13] present two methods that utilize an incremental support vector machine (SVM) along with clustering techniques to perform online HAR from streaming sensor data. They utilize a sliding sensor window that encodes the preceding sensor events in order to classify them. Other researchers have also experimented with similar sliding window techniques, ranging from static length sensor windows to dynamic time windows [14].

Drawbacks of traditional supervised HAR techniques are well summarised in [15], in which Ntalampiras and Roveri propose an interesting framework that utilizes Hidden Markov Models (HMM) for performing HAR and incremental dictionary learning to account for new user activities. The only assumption made in this framework is that the activities last for a specific amount of time. The researchers evaluated their technique on gyroscope and accelerometer data from activities such as walking, sitting and standing with promising results.

Most HAR approaches reviewed have utilized supervised learning or have been focused on simple physical activities. This paper instead proposes a semi-supervised framework for recognising and modelling more complex activities of daily living such as cooking, whilst also adapting to changing user behaviour.

3 Problem Formulation

There are two problems the research presented in this paper aims to solve – performing HAR from unlabelled data and tracking user activity in real-time. Traditional HAR approaches train a classifier on datasets which include a labelled dictionary of activities:

$\mathcal{A} = \{A_1, \dots, A_N\}$, where A_i denotes the i -th activity and N is the total number of activities. The classifier is then utilized to classify new occurrences of activities as one that is present in \mathcal{A} . This method has shown promising results, but does not account for any new activities that the user may start performing which were not present in the original dictionary or changes over time in the way the user performs those activities.

The second problem involves tracking the user activity which is crucial to providing real-time support. To formulate this problem, we first define the term user activity (or task) as a combination of multiple sets of user actions (or subtasks): $A = \{S_1, \dots, S_N\}$, where S_i denotes the i -th subtask and N is the total number of subtasks. The sequence in which these subtasks are performed may vary, as there may be multiple routes to the end goal of the task [16]. The challenge involved in tracking user activity is modelling the relationships between subtasks present in the activity, in order to predict the next sequence of subtasks.

Therefore, the goal of this research is to present a framework which can - automatically generate a dictionary of frequently repeated activities (defined in the next section) by analysing unlabelled sensor data; modify this dictionary according to changes in user behaviour; and model the relationships between the subtasks for each activity to allow for real-time tracking and support.

4 Adaptive Activity Recognition Framework

This section describes the overall framework of the proposed HAR system which aims to identify and model frequently repeated activities performed by the user in order to set up automated assistance/prompts and continually track changes in the performance of these activities for health assessment purposes. The system utilizes ambient and passive sensors embedded throughout the house rather than wearables or video data due to their greater acceptability as established in the literature review.

We first introduce a short-term system memory for storing and analysing sensor data. The system memory is specified as ' n ' days and is the number of preceding days's data the system would use from which to identify activities. If the present day is p , then a repeating activity will be identified as such if it is repeated at least once since day $p-n$. The value of n can be set to a larger integer for tracking longer-term activity model changes, but a smaller value of n favours a more adaptive system. A real-world implementation would require multiple non-overlapping system memories of different sizes to keep track of weekday and weekend behaviour separately (as these might be quite different). This paper only evaluates HAR of weekday behaviour.

The HAR system is divided into two parts, offline HAR and online HAR. Each of these parts are comprised of three steps (fig 1). The purpose of the offline HAR system is to discover activity clusters present in unlabelled user data at the end of each day and compare it to the activity clusters found in the previous n days to identify any repeating activities. This is then used to build a model of each identified repeating activity which can be used for performing online or real-time HAR for automated assistance/prompts and updating the activity model based on user actions.

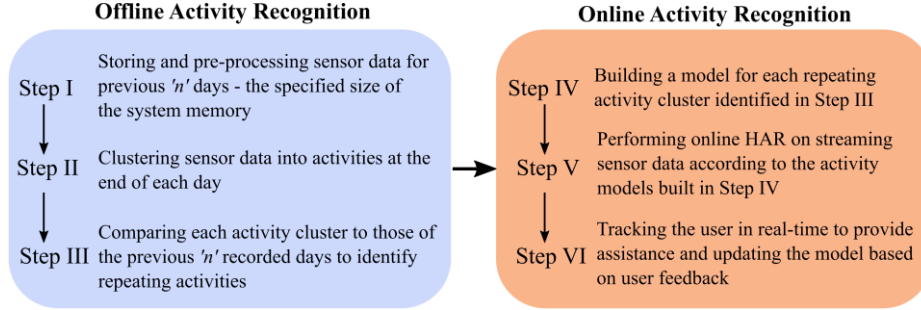


Fig. 1. Adaptive Human Activity Recognition Framework

These steps are repeated every day for the database of activity models to be created and updated. Any data older than the specified system memory is deleted, which results in old activities/habits being removed from the database. Any newly discovered repeating activities are added, while existing activity models are updated based on user behaviour (fig 2).

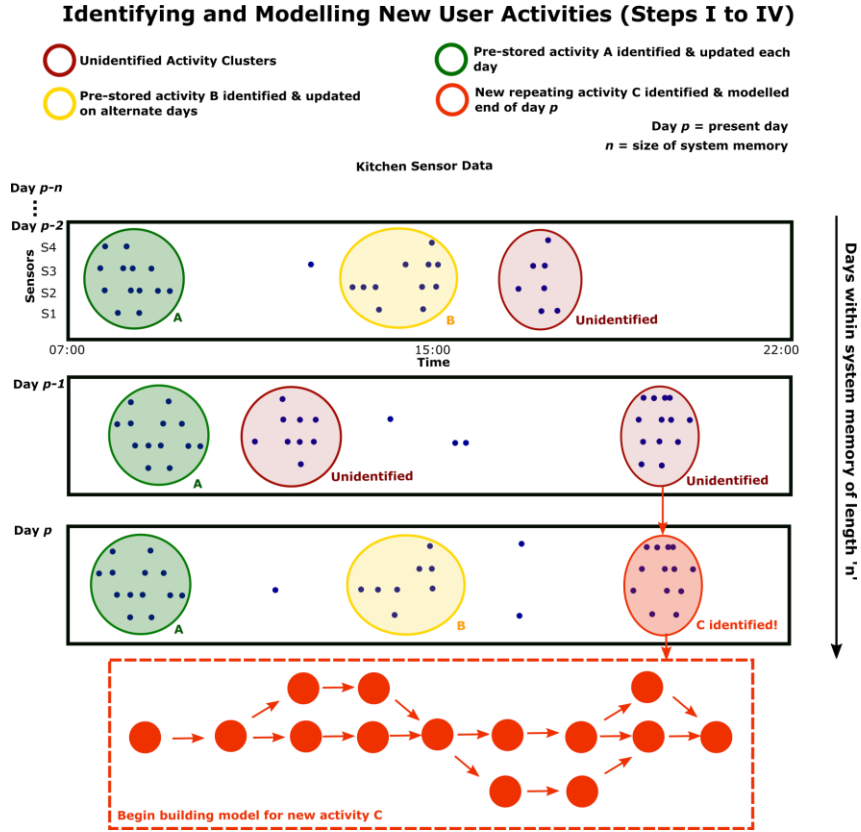


Fig. 2. Identifying and modelling a new repeating user activity introduced on day ' $p-1$ '

The following section describes the datasets used to develop and evaluate the system while Section 6 describes each step in more detail along with an explanation of the feature selection and the artificial intelligence implemented for each of the steps in this framework.

5 Aruba Dataset and Pilot Data Collection in the Assisted Living Studio

The Aruba CASAS study deployed multiple motion sensors in each room of an older female participant's house which yielded labelled activity data for a total of 7 months [5, 17]. Four weeks of kitchen data (two weeks from two different months) were selected at random from this dataset, split into two weeks each for training/developing and testing the offline HAR presented in this paper. Kitchen data was selected so that it could be augmented by additional complementary data collected in the Assisted Living Studio (ALSt) kitchen in the Bristol Robotics Laboratory. Weekend data was removed as only the weekday data was used in this study as explained in 3.1. While the presence of real-world noise in the Aruba dataset provided the ecological validity for testing the offline HAR component of this study, the sensor configuration lacked inclusion of other sensors such as contact sensors on drawers/cupboards, rendering the data insufficient to test the online HAR system presented in this paper. As such, a sample dataset was collected in the ALSt in the Bristol Robotics Laboratory (BRL) from participants recruited to perform typical kitchen-based activities, which served to augment the Aruba dataset. A Z-wave sensor network using openHAB was deployed in the ALSt which included Fibaro FGMS-001 motion sensors, Everspring SM810 magnetic contact sensors on cupboards, and TKB TZ69E wall plugs (to act as a sensor for the kettle). A total of six activity sessions were recorded from four participants on different days. Due to the nature of sensors used in the ALSt data collection, the sensor values only consisted of ON or OFF (0 or 1) values similar to the motion sensors deployed in the Aruba data collection. The number of ambient sensors deployed in the ALSt were also similar to the number of motion sensors present in the various rooms in the Aruba study to ensure compatibility between the two datasets.

The non-scripted kitchen activity selected for this study was a combination of two interleaved sub-activities - making a hot drink (such as tea, coffee or hot chocolate) and preparing a sandwich. The participants were provided with various items and ingredients appropriate for the selected activity and were asked to place these in the ALSt kitchen in a manner they were most comfortable with. They had access to two cupboards, one drawer, a counter-top, fridge, along with a kettle for boiling water and relevant utensils.

In the selected Aruba dataset there was a 30 to 60 minutes gap everyday between 16:00 and 17:00 after removing unlabelled activities, which was replaced by an ALSt participant's activity session data (fig 3). This was done to introduce real-world noise and data from the Aruba dataset, with the combined Aruba-ALSt data used for testing the offline HAR system, and the ALSt activity used for testing the online HAR system.

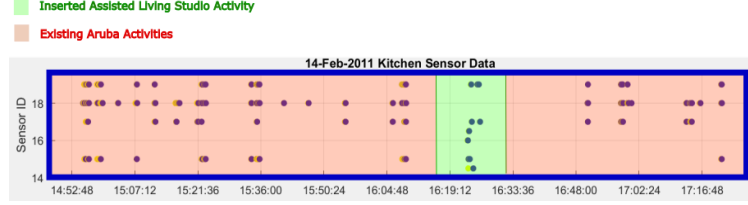


Fig. 3. ALSt activity inserted into Aruba CASAS dataset

6 Pilot Evaluation Study

This section describes an implementation of the adaptive HAR framework along with the data processing and machine learning techniques used. The weekday period was set to 5 days for pilot testing with the combined Aruba-ALSt data. Matlab was used for all data processing.

Step I: Data Pre-Processing

The sensor data obtained from the user for the specified system memory was separated by days and grouped by rooms –

Separation of Days. This was performed so that the sensor data for each day can be processed separately, and then compared to that of the previous day to identify repeating activities.

Grouping Sensors by Rooms. An assumption this paper made was that the activities are restricted by room, which would mean that if a single activity took place over three rooms (such as cleaning), it would be split into three activities (such as cleaning room 1, cleaning room 2, and cleaning room 3). The sensor data was separated and stored accordingly.

Step II: Discovering User Activities - Hierarchical Agglomerative Clustering

The purpose of this step was to discover user activities present in the data at the end of each day. This paper implements HAC to achieve this as used by [12], as it has been proved to be useful for discovering hierarchies of subtasks present within a task. The features used for clustering were the normalised sensor IDs and the timestamp at which the activity occurred. The maximum variance in the data was along the time axis, so HAC would cluster together data points that occurred within short time intervals. The distance calculated between data points was Euclidean. The linkage method was set to ‘single’, so that the shortest distance was used when clustering. A general rule for the cut off for HAC is to divide the square root of the number of data points by two [18]. After the analysis of the generated dendrograms from the training data, the cut-off formula was changed to

$$Cut\ off = \frac{\sqrt{no.of\ sensor\ events}}{1.5} \quad (1)$$

Therefore, the cut-off was set dynamically for each day depending on the total number of sensor events present.

Step III: Discovering Repeating Activities through Cluster Comparison

Once distinct activity clusters have been discovered, each activity cluster was compared to clusters from the previous five days to identify repeating activities. There were two features used for this –

1. *The time of the activity cluster.* Calculated as follows –

$$Activity\ time = \frac{time\ of\ last\ sensor\ event - time\ of\ first\ sensor\ event}{2} \quad (2)$$

This feature was used to compare the time at which the activity occurred on different days. This means that an activity had to be repeated within the same time frame each day for it to be considered a repeating activity (set as one hour in this study after training data analysis). This is a shortcoming of this current implementation of the Adaptive Activity Recognition Framework which could be overcome by a future implementation of dynamic time warping for cluster comparison [19].

2. *The total number of sensor events present in the activity cluster.* This feature was used to represent the sizes of the activity clusters for comparison. In this paper clusters greater than 0.75 and smaller than 1.67 of the original were considered similar (thresholds set after analysis of training data).

Sensor data belonging to activity clusters identified as similar/repeating were stored using a unique identifier to refer to the activity. The data from each identifier was then used to build the activity model for that identifier in the next step.

Step IV: Building User Activity Model using Bayesian Networks.

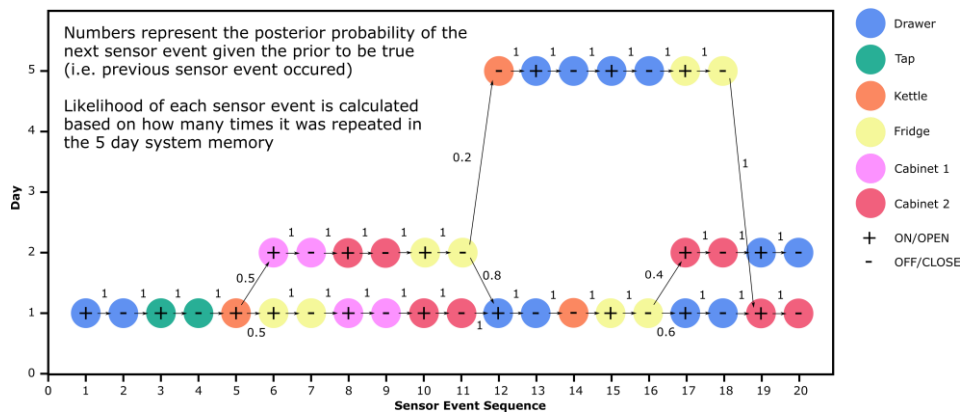


Fig. 4. Bayesian network for Assisted Living Studio Activity for Participant B

Sensor data from previous days' of user activity were used to generate a Bayesian net of the activity (fig 4). This involves identifying and collapsing repeating sequences of sensor events from each previous day. Fig 4 encodes all possible sequences of events the user has performed for the activity in the last five days. Days 3, 4 and 6 are empty as the activity from those days can be reproduced using the sequence of events from days 1, 2 and 5. The probability of the user performing the next sequence of events depends on how many times the user has repeated the same sequence in the past.

Step V: Real-Time HAR Using Sliding Sensor Window

Once the Bayesian model is generated, a sliding window comprising of three sensor events is used for performing real-time HAR, similar to [13] as shown in fig 5 (Note: this is different from the system memory described in section 3.1). This is achieved by first locating the unique activity models from previous days present in the database (generated in step IV) which occurred within an hour of present time. Then comparing the most recent three sensor events (sliding window size = 3) to the first three sensor events present in the activity models to identify the activity being performed.

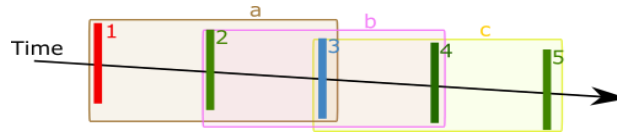


Fig. 5. Sliding window for sensor events – numbered lines are sensor events, lettered boxes are sliding windows of 3 sensor events.

Step VI: Updating User Activity Model through User Feedback

Following real-time recognition, the SH algorithm needs to track the user in to order detect any change in behaviour and automate prompts accordingly. For this, a Matlab script was created to use the Bayesian nets created in the previous step to predict the next most probable move by the user. If the user misses a step, the script keeps track of the user by jumping ahead to reach the step the user is at. However, the script also keeps track of any missed steps in order to provide a prompt to the user to remind him of the missed step. According to the user's response of either correcting for the missed step by doing specific missing action, responding to the prompt with text input, or ignoring the prompt, the Bayesian net for the activity is updated accordingly.

7 Results and Discussion

HAC (Step II) was evaluated on 12 days of Aruba weekday data for kitchen activities between the dates 3rd January 2011 to 18th February 2011 with weekend data removed, as this study only evaluated the weekday system memory. Recognised activities were separated into activities identified precisely and activities identified with noise. An activity was classified as identified precisely if the number of sensor events in the identi-

fied activity cluster were same as the number of events listed in the dataset. Some activity clusters contained additional sensor events not part of the label which were classified as “identified with noise” (fig 6). The accuracy was calculated as follows

$$Accuracy = \frac{\text{Number of activities identified precisely}}{\text{Total number of activities labelled in dataset}} \times 100 \quad (3)$$

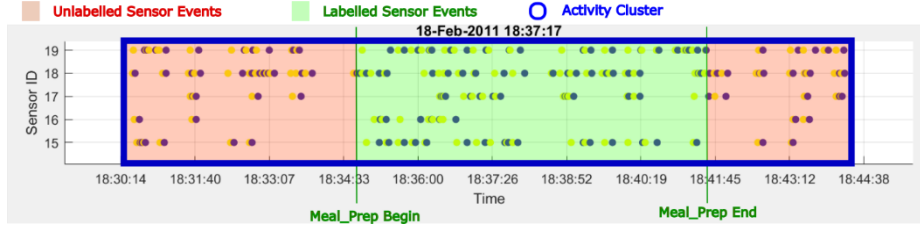


Fig. 6. “Identified with Noise” Meal Preparation Activity, 18 Feb 2011

The overall accuracy of the HAC was 71% for the Aruba dataset in separating sensor events into separate activity clusters. The system identified six separate “Meal Prep” clusters as there were a maximum of six separate instances of the “Meal Prep” activity present in the same day, which could signify different variations of the activity such as preparing different types of meals, making a hot drink, etc.

Cluster comparison (Step III) was then performed to identify the same activities repeating on different days. The system compared the six “Meal Prep” clusters from step II to identify repeating instances of the activity present on other days with an F1-score of 0.93, while the score for the identifying repeating “Wash Dishes” activity was 0.35. For comparison, authors in [13] used the Aruba dataset to achieve an F1-score between 0.6 and 0.7 for “Meal Prep” activities and 0 for “Wash Dishes”. They concluded that the classifier failed to identify the “Wash Dishes” activity due to imbalance in the dataset because of the presence of unlabelled sensor events that dominates it. This also seems to be the case here as the “Wash Dishes” activity was performed less consistently than “Meal Prep” in the selected data and the system often classified unlabelled data as “Wash Dishes” during cluster comparison, resulting in the low F1-score. It must also be noted that authors in [13] used six weeks of Aruba dataset including training and testing, while this paper only used 12 days for both, making the former’s approach more robustly tested. The system also identified instances of unlabelled data as an activity which were removed from the results as they cannot be verified due to missing labels (fig 7).

The F1-score for the cluster comparison of the ALSt activity from the combined Aruba-ALSt dataset was 0.96. Real-time activity recognition (Step V) for the ALSt participant activity from day 2 onwards occurred at the beginning of the activity (3rd sensor event) 64% of the time, and near the midway point (5th to 7th sensor event) 96% of the time. On day 6, real-time activity recognition occurred at the beginning of the activity 100% of the time for all participants. This is because the Bayesian activity model improves each day utilizing the previous 5 days of data. It must also be noted

that the earliest real-time detection can only happen at the 3rd sensor event due to the size of the sliding sensor window being three sensor events.

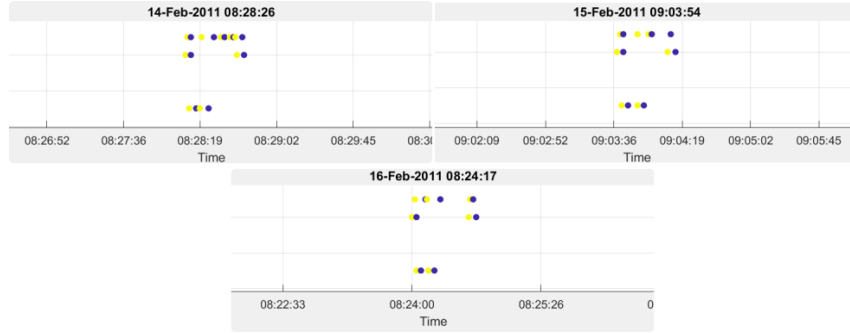


Fig. 7. An unlabelled repeating activity cluster identified and isolated

The last experiment was to evaluate how well the system can track the user to detect missed steps using the Bayesian network and provide prompts (Step VI), which was achieved by simulating the user forgetting to open/close the cabinet and the fridge during the activity on day 5 and day 6. The system detected these missed steps in 75% of the instances on day 5 and 87.5% of the instances on day 6.

8 Conclusions and Future Work

A significant problem with current HAR techniques is the requirement of large amounts of expertly labelled training data for each specific end-user. This paper presented an alternative approach to HAR which utilises a short-term system memory with specified decay length to recognise repeating activities from unlabelled data and track them in real-time. The user activity models generated by this system can be utilized to offer additional support, for example, if the user forgets or misses a step in a known activity pattern such as forgetting to turn off the hob while preparing a meal, the system can then remind them through automated prompts. Additionally, this approach can be used to track changes in user activities which could be an indication of deteriorating health conditions and alert the carer of significant changes based on a personalised threshold. An implementation of the proposed HAR framework was also pilot tested using the Aruba dataset augmented with sample data collected in the ALSt in the Bristol Robotics Laboratory with promising results. Future work will include testing this framework on different datasets and attempting to improve the cluster comparison performance by implementing techniques such as dynamic time warping [19].

References

1. Demiris, G., Hensel, B.K.: Technologies for an aging society: a systematic review of “smart home” applications. *Yearb Med Inf.* 3, 33–40 (2008).
2. Kvedar, J., Coye, M.J., Everett, W.: Connected health: a review of technologies and strategies to improve patient care with telemedicine and telehealth. *Health Aff.*

- (Millwood). 33, 194–9 (2014).
3. Shoaib, M., Bosch, S., Incel, O., Scholten, H., Havinga, P.: A Survey of Online Activity Recognition Using Mobile Phones. *Sensors*. 15, 2059–2085 (2015).
 4. Ranasinghe, S., Al MacHot, F., Mayr, H.C.: A review on applications of activity recognition systems with regard to performance and evaluation. *Int. J. Distrib. Sens. Networks*. 12, 1550147716665520 (2016).
 5. CASAS: Smart Home Projects, Washington State University (USA), <http://casas.wsu.edu/>.
 6. Pal, S., Feng, T., Abhayaratne, C.: Real-time recognition of activity levels for ambient assisted living. 5th IEEE Int. Conf. Consum. Electron. - Berlin, ICCE-Berlin 2015. 485–488 (2016).
 7. Roggen, D., Förster, K., Calatroni, A., Tröster, G.: The adARC pattern analysis architecture for adaptive human activity recognition systems. *J. Ambient Intell. Humaniz. Comput.* 4, 169–186 (2013).
 8. Cardinaux, F., Bhowmik, D., Abhayaratne, C., Hawley, M.S.: Video based technology for ambient assisted living: A review of the literature. *J. Ambient Intell. Smart Environ.* 3, 253–269 (2011).
 9. Basu, D., Moretti, G., Sen Gupta, G., Marsland, S.: Wireless sensor network based smart home: Sensor selection, deployment and monitoring. In: 2013 IEEE Sensors Applications Symposium, SAS 2013 - Proceedings. pp. 49–54 (2013).
 10. Rashidi, P., Mihailidis, A.: A survey on ambient-assisted living tools for older adults. *IEEE J. Biomed. Heal. Informatics*. 17, 579–590 (2013).
 11. Boise, L., Wild, K., Mattek, N., Ruhl, M., Dodge, H.H., Kaye, J.: Willingness of older adults to share data and privacy concerns after exposure to unobtrusive home monitoring. *Gerontechnology*. 11, 428–435 (2013).
 12. Kwon, Y., Kang, K., Bae, C.: Unsupervised learning for human activity recognition using smartphone sensors. *Expert Syst. Appl.* 41, 6067–6074 (2014).
 13. Yala, N., Fergani, B., Fleury, A.: Feature extraction and incremental learning to improve activity recognition on streaming data. In: 2015 IEEE International Conference on Evolving and Adaptive Intelligent Systems, EAIS 2015. pp. 1–8 (2015).
 14. Krishnan, N.C., Cook, D.J.: Activity recognition on streaming sensor data. *Pervasive Mob. Comput.* 10, 138–154 (2014).
 15. Ntalampiras, S., Roveri, M.: An Incremental Learning Mechanism for Human Activity Recognition. 2016 IEEE Symp. Ser. Comput. Intell. 1–6 (2016).
 16. Becker, M., Prueckner, S., Storf, H., Kleinberger, T., Schmitt, M., Bomarius, F.: An Event-Driven Approach to Activity Recognition in Ambient Assisted Living. *LNCS*. 5859, 123–132 (2009).
 17. Cook, D.J.: Learning Setting- Generalized Activity Models for Smart Spaces. *IEEE Intell. Syst.* 27, 32–38 (2012).
 18. Richard Duda Peter Hart, D.S.: *Pattern Classification* (Second Edition). (2001).
 19. Petitjean, F., Ketterlin, A., Gançarski, P.: A global averaging method for dynamic time warping, with applications to clustering. *Pattern Recognit.* 44, 678–693 (2011).