

Monitoring and Modeling Simple Everyday Activities of the Elderly at Home

George Papamathaiakis, George C. Polyzos, George Xylomenos

Mobile Multimedia Laboratory

Department of Informatics

Athens University of Economics and Business

Athens 104 34, Greece

{papamath06, polyzos, xgeorge}@aueb.gr

Abstract—We present our work on a sensor-based smart system that is automatically trained to recognize the activities of individuals in their home. In this paper we present and analyze a method for recognizing indoor the everyday activities of a monitored individual. This method is based on the data mining technique of association rules and Allen’s temporal relations. Our experimental results show that for many (but not all) activities, this method produces a recognition accuracy of nearly 100%, in contrast to other methods based on data mining classifiers. The proposed method is accurate, very flexible and adaptable to a dynamic environment such as the “Smart Home” and we believe that it deserves further attention.

I. INTRODUCTION

The world population of people over the age of 65 is growing rapidly at a rate of 800,000 per month. As a consequence, the healthcare and elderly-care market already constitutes a major part of the economy and it will only expand in years to come. Recent advances in sensor technology, cellular networks and information technology promise to improve the well-being of the elderly by assisting them in their everyday activities and monitoring their health status, thus enabling them to lead their lives to a larger extent independently from healthcare institutions and caretakers.

However, the comparatively slow adoption of such systems indicates that there are certain factors prohibiting their acceptance. For example, most home-care systems monitor the health of individuals suffering from diagnosed chronic diseases (heart disease, lung disorders, diabetes, etc.); as such, they depend on customized, costly, and difficult to use equipment, such as cardio graphic monitors, often limiting patient mobility. Less attention has been given to monitoring and maintaining the personal wellness of elderly people who *have not been diagnosed with any serious or chronic disease* and who, therefore, wish or should be encouraged to live a normal life. The quality of life of this population depends upon, and may decline as a result of, combined and (in some cases) not easily measurable physical and psychological factors. Existing work in improving their lives is limited to interactive facilities for consultation with doctors, which, while simplifying regular examinations, do not take advantage of

the potential offered by the latest advances in information and communication technologies.

While sudden changes and, more generally, unexpected patterns in the everyday activities of the elderly are difficult to detect via medical data only, it is very likely that i) such unexpected patterns can be detected by combining machine learning and input from a larger variety of sensors, and that ii) they can be used as signs that a certain disease, illness or health condition is getting worse, or is about to occur. Of course, detecting these early signs is far from trivial, especially if we consider the uniqueness of each individual’s personality, routine, and the diverse effects of different diseases upon this routine. Furthermore, the quest for early detection should not come at the cost of confining the elderly to their homes or limiting their everyday activities, nor should it rely on the assumption that the elderly will have to learn to control complex equipment.

Unfortunately, researchers do not currently have the tools to recognize human activities using a set of simple, easy to install ubiquitous sensors, nor do they understand what modifications are necessary in pattern recognition algorithms to achieve this. But, if it is possible to develop systems that recognize an individual’s everyday activities, researchers may be able to automatically detect changes in the behavioral patterns of people at home that indicate declining health [4].

The aim of our work is to create a smart system that will be adaptable to individuals, will be able to recognize their activities and will help their well being by raising alarms when a potential departure from routine or desirable behavior is detected (for example, the individual did not eat lunch or appears to not take their medication at prescribed times). In other words, the smart system will be trained to recognize the activities of an individual inside his or her personal space. With this approach, the monitored individual plays a critical role in the system’s training as the system learns to distinguish the everyday activities and habits of each person.

In this paper, we present and study a method of indoor everyday activity recognition based on the data mining technique of association rules and Allen’s temporal relations [4]. We also present a simple activity recognition method based on the C4.5 classification algorithm [8], [9] and compare it with our method based on association rules.

This research has been supported in part by project ARCHANGEL through a Cell Phone as a Platform for Healthcare award from Microsoft Research.

II. DATA SETS

According to Szalai [1], an individual's actions can be characterized and grouped in certain categories based on her everyday habits. The analysis of data acquired by recording the everyday activities of individuals using "tape on and forget" sensors in private spaces, presented in [2], [3], shows that we can detect and identify certain activities of the monitored individual with an accuracy of up to 89%. The identified activities can be processed to produce a useful image of the everyday habits of the monitored persons. The results presented in [2] were based on data obtained from the everyday activities of two individuals, 30 and 80-year-old respectively. The set of actions recorded in their private spaces was obtained by sensors positioned in nearly 80 different locations inside their apartments. These actions were then grouped in categories, which helped make the processing of the samples easier, with the final aim being the extraction of in-home patterns of common activities. This data set has been made public by the research team and is therefore used in our study as an important input for the custom in-home models that we will produce to test the system's sensor-based algorithm.

III. DEFINITIONS

A. Association Rules

An association rule finds relationships among large sets of data items. Association rules indicate the attribute value conditions that occur frequently together in a given dataset. In association analysis, the antecedent and consequent of a rule are two sets of items (called itemsets) that are disjoint.

An association rule includes two numbers that express the degree of uncertainty about the rule. The first number is called the *support* for the rule. The support is simply the number of transactions that include all items in the antecedent and consequent parts of the rule; the support is sometimes expressed as a percentage of the total number of records in the database. The second number is known as the *confidence* of the rule. The confidence is the ratio of the number of transactions that include all items in the consequent as well as the antecedent (i.e., the support) to the number of transactions including all items in the antecedent only.

B. Temporal Relations

Activities in a smart environment include physical activities as well as interactions (with objects). For example, activities may include walking, resting on a couch and using the coffee-machine. An important observation made in [4], upon which we base our approach, is that these activities are not instantaneous, but have distinct start and end times. In addition, well-defined time relations exist between the events constituting an activity. These time relations can be represented using Allen's temporal relations. Allen proposed describing activity scenarios with time relations and presented a set of temporal relations between events (see Figure 1 for some examples) [5]. These temporal relations are important in the determination of the monitored user's indoor activities and can be used for knowledge and pattern discovery in day-to-day activities.

Temporal Relations	Representation	Interval Constraints
X AFTER Y		$Start(X) > Start(Y)$ $End(Y) < Start(X)$
X DURING Y		$Start(X) > Start(Y)$ $End(X) < End(Y)$
X MET-BY Y		$Start(X) = End(Y)$
X STARTS Y		$Start(X) = Start(Y)$ $End(X) \neq End(Y)$
X STARTED-BY Y		$Start(Y) = Start(X)$ $End(X) \neq End(Y)$
X FINISHES Y		$Start(X) \neq Start(Y)$ $End(X) = End(Y)$
X FINISHED-BY Y		$Start(X) \neq Start(Y)$ $End(X) = End(Y)$
X EQUALS Y		$Start(X) = Start(Y)$ $End(X) = End(Y)$
X OVERLAPPED-BY Y		$Start(Y) < Start(X)$ $Start(X) < End(Y)$ $End(Y) < End(X)$

Fig. 1. The most important temporal relations used for our method.

IV. ACTIVITY RECOGNITION

A. Detecting Temporal Relations

The first step in activity recognition is to process the raw data to find the temporal relations between sensor events. This is achieved by using a simple algorithm which takes the timestamp of each event that occurred, using the available data [2], and identifying the temporal relationships between pairs of consecutive events based on the constraints formulated in Figure 1. The pseudo code for the temporal analyzer tool is described in the following algorithm.

```

Input: sensor events, start time, end time of events
Repeat While [∃ Unprocessed Event]
  ⇒ Find Start & End time of next events
  ⇒ Compare Start & End times of past and next event
  ⇒ Identify relation type between event pair from the possible
    relation types (see Figure 1)
  ⇒ Record those 2 sensors (IDs) & create the temporal relation
  ⇒ Increment Event Pointer (go to the next event)
Loop Until end of Events.
  
```

For example, assume that we are given as input the following consecutive events:

Date	Start	End	Sensor	Object	Room
109	13624	19148	137	5	2
109	15939	18176	109	5	9
			(i.e.)		
			↓		
19/4/2003	3:47:04	1:32:04	137	'Door'	'Bathroom'
19/4/2003	4:25:39	0:37:17	109	'Light'	'Bathroom'

The output of the algorithm will be the following temporal relation between sensors 137 and 109:

"109 DURING 137"

The key premise of our work is that each activity can be described with such a set of temporal relations. In particular, we have noticed that similar activities lead to the appearance of a common set of temporal relations between specific sensors. In other words, if we can isolate the most important temporal relations characterizing activities, it would be easy to describe any activity with high prediction accuracy. The solution to this problem is based on data mining techniques and the association rules.

B. Mining Association Rules

In order to create association rules we use the Weka tool [7]. Using the Apriori algorithm [6] of this tool, we can identify the most important temporal relations associated with each activity. More specifically, we first provide as input to the Weka tool the temporal relations which describe a specific activity in the following format:

SensorID_X - SensorID_Y,	Relation
S_141-S_98,	DURING
S_137-S_141,	AFTER
S_141-S_137,	DURING
S_110-S_141,	AFTER
S_137-S_110,	DURING
S_112-S_137,	OVERLAPBY
S_106-S_112,	DURING

Based on this set of input temporal relations, the Weka tool outputs the most important association rules for the activity. In particular, the output of the Weka tool for a particular activity has the following format:

1. S_141-S_141	29	⇒	Relation=AFTER	29	conf:(1)
2. S_101-S_101	15	⇒	Relation=AFTER	15	conf:(1)
3. S_115-S_115	4	⇒	Relation=AFTER	4	conf:(1)
4. S_100-S_141	3	⇒	Relation=AFTER	3	conf:(1)
5. S_100-S_100	3	⇒	Relation=AFTER	3	conf:(1)
6. S_137-S_137	5	⇒	Relation=AFTER	4	conf:(0.8)

For example, the events in rule 1 appeared 29 times in temporal relations in this activity, in all cases the AFTER relation, therefore the confidence of this association rule is $29/29 = 1$. In contrast, for rule 6 we had 5 appearances of the events only 4 of which were in the AFTER relation, therefore this association rule has a confidence of $4/5 = 0.8$. From the association rules in the Weka output, we choose those rules providing the desired support (**support** \geq **minsup**) and confidence (**confidence** \geq **minconf**) as the ones that can best characterize the input activity.

C. Detecting Activities

Using to the procedure described above, we discover the most important temporal relations between successive pairs of sensors. We can then use a simple algorithm to detect the type of the activity represented by a sequence of temporal relations between successive events by calculating the *importance degree* for each activity. The importance degree for an activity is the sum of confidence percentages for the association rules of the activity detected in the data set.

For example, assume an activity that can be described based on the association rules of the previous example above. If

Rule 1 and **Rule 6** appear in the input 3 times and 2 times, respectively, then the activity's *importance degree* will be:

$$1 \times 3 + 0.8 \times 2 = 4.6$$

The activity with the highest importance degree is the most preponderant and, probably, the one actually performed.

V. EXPERIMENTS AND RESULTS

For the experimental evaluation of our association rules technique we use the following scheme. Using a set of events for each activity, we initially produce the temporal relations between successive events and find the most important temporal relations, based on the association rules. Then, using the most important temporal relations for each activity, we try to recognize the activities that were actually executed. Activity recognition is always based on the *importance degree* which characterizes each activity.

As discussed above, the available data used in our work [2] consist of activities that were recorded in two different apartments for different people. Thus, we executed our experiments for each one of the individuals and present separately the activity recognition results. It should be noted that the monitored individuals themselves indicated the activities they were performing during the monitoring process out of a given set of activities, therefore we have a concrete benchmark against which to test our scheme.

It is important to notice that all these experiments were executed using *variable* lower bounds of support (**support** \geq **min-sup**), because the number of events in each activity sometimes differed from one another. For example, let us assume an activity with temporal relations that are reused a maximum of three (3) times. If the lower bound of support is five (5) for all the activities, it is impossible for the Apriori algorithm [6] to output temporal relations that are reused less than five (5) times. Therefore, depending on the number of events in each activity, we use an appropriate lower bound for association rule support. This adaptation seems to be the key to the success of our method.

A. 1st Data Set (Home #1)

We initially train the system with all available data and then try to discover all the activities in the same data set. In this manner we understand if the system is well trained or not and if it can recognize the dataset that it trained for. In Table I we present the recognition precision for the most important activities. The accuracy of the method is the ratio of achieved activity recognitions to the total number of activity appearances. We consider as the most important activities those with a large number of occurrences in the data set, as these are more likely to lead to concrete conclusions.

According to these results, our method achieves good enough detection accuracy, as it approaches 100% in most cases. However, the method does not give equally good results for the **Toileting** activity. This is the case because of the instability of our data samples. In a large part of these traces, the events that constitute the **Toileting** activity follow a concrete pattern. This fact strengthens these events' place in

TABLE I

RECOGNITION OF ACTIVITIES – 2 WEEKS (TRAINING FOR 2 WEEKS)

Accuracy	Activity
0.85	Bathing
0.88	Grooming
0.92	Doing laundry
1.0	Preparing lunch
0.5	Toileting

the most important rules, because of their appearance many times in the samples. Nevertheless, the remaining events of this activity are less frequent, thus they do not affect the association rules. As a result, while half of the events of an activity achieve a very good accuracy rate, the remaining events achieve a poor accuracy rate. This problem reduces the total activity recognition accuracy of our system due to the **Toileting** activity. A solution for this problem is to also consider less important association rules, something which decreases the value of our method however, since we want to recognize activities based on the important association rules.

TABLE II

RECOGNITION OF ACTIVITIES – 1 WEEK (TRAINING FOR 1 WEEK)

Accuracy	Activity
0.72	Bathing
0.53	Grooming
1.00	Doing laundry
0.72	Preparing lunch
0.86	Toileting

Some researchers believe that one week is a sufficient period for training a system for a smart environment. For this reason, we also trained our system based on the data from the first week and evaluated the system's recognition accuracy based on the data from the second week. This time, our results were not expected to be as spectacular as before, as we tried to discover activities with events that are probably unknown, as they do not appear in the training data – they only appear during the second week. However, the results shown in Table II, contrary to our expectations, indicate that the recognition accuracy is actually quite good.

B. 2nd Data Set (Home # 2)

We next examine the quality of the method and its precision in activity recognition over the data from the second home. Again, we first train the system using all available data and then try to discover all the activities from the same data set. In Table III we present the recognition precision of our method for the most important activities.

TABLE III

RECOGNITION OF ACTIVITIES – 2 WEEKS (TRAINING FOR 2 WEEKS)

Accuracy	Activity
0.93	Preparing breakfast
0.53	Preparing lunch
0.55	Listening to music
0.91	Toileting
0.75	Watching TV

Based on these results, we conclude that in the second home the method behaves overall satisfactorily. However, this is not the case for all the activities. In Table III, the activities **Preparing Lunch** and **Listening to Music** barely exceed an accuracy rate of 50%. An explanation of this problem is that these two activities contain events which also appear in other activities. For example, during the activity **Listening to Music** the monitored individual may simultaneously be preparing something to eat or visiting the bathroom. Even if these events (actions) are parts of the e.g. **Preparing Breakfast** or **Toileting** activities, the main activity is **Listening to Music**. This is perhaps the major problem of our data set and unfortunately, it is not easy to overcome. The main cause of the "erroneous" event classification is the way that the events were originally classified by the monitored persons themselves. In addition, the **Preparing Lunch** activity contains some important events, which are also included in the **Preparing Breakfast** activity. Since the **Preparing Breakfast** activity will probably be marked with a higher importance degree, the system will not give the **Preparing Lunch** activity such a high importance. A solution for this problem is a more careful and intelligent retrieval of association rules for activities.

TABLE IV

RECOGNITION OF ACTIVITIES – 1 WEEK (TRAINING FOR 1 WEEK)

Accuracy	Activity
0.10	Preparing breakfast
0.64	Preparing lunch
1.00	Listening to music
1.00	Toileting
1.00	Watching TV

The next experiment is again based on training the system with the first week's data and evaluating its prediction ability using the second week's data. In this experiment, we expect very good results, because most of the important events for the first week constitute the most important information for the complete data set. The results obtained from that experiment (seen in Table IV) verify our predictions by presenting high event recognition accuracy. Nevertheless, we observe an inversion in the accuracy for the **Preparing Breakfast** and **Preparing Lunch** activities. Contrary to the previous experiment, the **Preparing Breakfast** activity results in an excessively low percentage (10%) in contrast to the high increase of the percentage for the **Preparing Lunch** (64%) activity. As before, **Preparing Lunch** and **Preparing Breakfast** are related activities difficult to distinguish. This observation becomes very obvious with the results of this experiment, where the accuracy reduction for one activity results in an accuracy increase for the other. This is due to our choice to emphasize the most important temporal relations between events (via the association rules). Note however that if we were to group the similar activities **Preparing Lunch** and **Preparing Breakfast** into a single activity (perhaps **Preparing Meal**) we would end-up with a recognition accuracy of 85% for the grouped activity.

VI. CLASSIFICATION METHOD

An alternative method for activity recognition is to take into account the combination between the sensors that appear in the activities and the number of their occurrences in each activity. For example, the TV sensor is activated 4 times in the activity **Watching TV**. For this particular scheme, we consider appropriate data mining techniques and employ the corresponding classification technique.

We used again the Weka tool [7] for classification purposes. We organized the input data for Weka in a file of records defining for each activity the number of times that a sensor was activated, as described in Table V. The first line shows the sensors participating in the activity, while the second line shows the number of occurrences for each sensor in the corresponding activity.

TABLE V
INPUT FORMAT FOR CLASSIFICATION SCHEME

Sensor	51	...	75	...	141	Activity
# Events	2,	...	3,	...	5	Watching TV

After organizing the data for all the activities using the aforementioned procedure, we applied as the classification scheme the **J48** algorithm supported by Weka [10] (classification algorithm C4.5 [8], [9]). This algorithm (J48) produced the best results in our experiments among various other classification algorithms.

VII. CLASSIFICATION VS. ASSOCIATION RULES

In general, activity recognition using association rules produced better results than classification. Table VI supports this statement for the first dataset (Home #1). Note that for the classification scheme we used 60% of the dataset for training, i.e. more than one week, and 40% for testing the accuracy of the algorithm, as this provided the best accuracy.

TABLE VI
ACCURACY FOR EACH METHOD (HOME #1)

Association Rules	Classification J48	Activity
0.72	1.00	Bathing
0.53	0.158	Grooming
1.00	0.80	Doing laundry
0.72	0.00	Preparing lunch
0.86	0.824	Toileting

As we can see from Table VI, association rules are superior compared to the classification method. The latter is only better in the **Bathing** activity, but even for this activity the association rules provided a satisfactory result. We should also add that the result of zero success for the classification method in recognizing the **Preparing Lunch** activity is due to the activities distribution of our data: there are not enough **Preparing Lunch** activities into the second 40% of the dataset for the classification experiment.

Association rules also produced better results for the second dataset (Home #2), as depicted in Table VII. Note that for the classification scheme we used 70% of the dataset for training, and 30% for testing the accuracy of the algorithm, as this

TABLE VII
ACCURACY FOR EACH METHOD (HOME #2)

Association Rules	Classification J48	Activity
0.10	0.667	Preparing breakfast
0.64	0.571	Preparing lunch
1.00	0.667	Listening to music
1.00	0.909	Toileting
1.00	0.60	Watching TV

provided the best accuracy. We observe that the only activity where problems arise for the association rules is **Preparing Breakfast**. This situation is again caused by our data set (i.e. **Preparing Lunch** and **Preparing Breakfast** are related activities difficult to distinguish) and demands further special treatment in order to be addressed.

VIII. CONCLUSIONS

The experimental results presented in this paper support our claim that our method based on association rules constitutes a very good choice in the field of activity recognition. For some basic activities this method produces a recognition accuracy of nearly 100%, in contrast to classification methods [2] based on data mining classifiers. These classification methods gave a maximum recognition accuracy of 89% via the use of simple classifiers such as the Naïve Bayes classifier. In this paper we also considered another simple method for activity recognition, based on the C4.5 classifier [8], [9] (J48 in the Weka Toolkit). This classification technique provided good results, but not as good as the method based on association rules.

For certain activities, however, the association rule method met with a low rate of success. Since we do have reasonable explanations for these results, we believe that the proposed method is appropriate and one can further improve the accuracy of the system with additional manipulations.

In conclusion, the association rules method is more accurate, flexible and adaptable in a dynamic environment such as the "Smart Home." Therefore, it deserves further attention and related research in this direction is needed.

REFERENCES

- [1] A. Szalai, "The use of time. Daily activities of urban and suburban populations in twelve countries." *The Economic Journal*, 84:335, 1974.
- [2] E. M. Tapia, S. S. Intille, and K. Larson, "Activity recognition in the home setting using simple and ubiquitous sensors," Proc. Pervasives 2004, Vienna, Austria, 2004.
- [3] <http://courses.media.mit.edu/2004fall/mas622j/04.projects/home>
- [4] V.R. Jakkula, and D.J.Cook, "Mining Sensor Data in Smart Environment for Temporal Activity Prediction," Poster session at the ACM SIGKDD, San Jose, CA, August 2007.
- [5] J.F. Allen, and G. Ferguson, "Actions and Events in Interval Temporal Logic," Technical Report 521, July 1994.
- [6] V.R. Jakkula, and D.J. Cook, "Learning Temporal Relations in Smart Home Data," Proc. Second International Conference on Technology and Aging, Canada, June 2007.
- [7] <http://www.cs.waikato.ac.nz/ml/weka/>
- [8] <http://www.cis.temple.edu/~ingargio/cis587/readings/id3-c45.html>
- [9] http://en.wikipedia.org/wiki/C4.5_algorithm
- [10] http://grb.mnsu.edu/grbts/doc/manual/J48_Decision_Trees.html