TOWARDS THE DETECTION OF TEMPORAL BEHAVIOURAL PATTERNS IN INTELLIGENT ENVIRONMENTS

F. RIVERA-ILLINGWORTH, V. CALLAGHAN, H. HAGRAS
UNIVERSITY OF ESSEX, UNITED KINGDOM

ABSTRACT

Ubiquitous computing applications propose new and creative solutions to every day needs. This paper addresses the issue of recognition of every day activities inside pervasive domestic environments in order to identify patterns of behaviour that can be later used to support care systems by detecting changes to those patterns. Our system uses a temporal neural-network-driven embedded agent able to work with online, real-time data from unobtrusive low-level sensors and actuators. We present experimental results that show our agent is able to detect temporal patterns along with spatial similarity associations found in human behaviours and activities, in everyday living environments.

Keywords: Behaviour detection, Intelligent agents, Abnormality detection, Intelligent environments, Temporal patterns

1. INTRODUCTION

New information paradigms such as pervasive computing and ambient intelligence propose innovative and creative solutions to every day needs.

Intelligent environments are closely related with the term ubiquitous or pervasive computing. This term was originated by Mark Weiser (1), who had the vision of "people and environments augmented with computational resources that provide information and services when and where desired". An intelligent environment has been defined by Zita and Yanco (17) as (e.g. a room, a house, an office, etc) as a place that "has sensors and actuators that monitor the occupants, communicates with each other, and intelligently supports the occupants in their daily activities". These environments can simply be automated systems or more complex adaptive systems that use some sort of learning, reasoning or planning to adapt themselves to the inhabitant's behaviour.

The potential of these environments has been demonstrated in various domains such as offices, classrooms, energy saving/management, consumer satisfaction, supportive environments, etc. One of the more important and useful metrics used to study human behaviours inside a domestic environment are those by Mihailidis et al. (9), Munguia et al (12), Mynatt and Rogers (13), Philipose et al (14), Wilson and Atkenson (16) and, Zita and Yanco (17) which track the Activities of Daily Living (ADLs) and the Instrumental Activities of Daily Living (IADLs).

The research presented in this work addresses the issue of recognition of every day activities inside pervasive domestic environments in order to identify patterns of behaviour that can be later used to support care systems by detecting alterations to those patterns.

The recognition of human activities poses several challenges due to the diverse number of ways people perform those activities, the configuration of the sensory system inside the environment, and the architecture used to detect and classify the activities.

One area that has received much attention is the sensory system used for the activity detection. It has been recognized that monitoring techniques that are relatively automated and unobtrusive are much more likely to be successful. The consequence of those techniques according to Holly et al (6) is that the acquired data is going to be noisier and requires more sophisticated algorithms for inferring the current state, but the data will be more continuous and not dependent on people adherence for success.

Beaudin et al (2) mention that some desirable characteristics for the sensor set are that it should be unobtrusive and that no modifications, or only minor modifications, to the environment should be needed in order to deploy them. The sensors also need to be reliable, require no maintenance and ideally be cheap, so they can be deployed in large quantities.

The use of very simple on/off sensors such as motion detectors, pressure sensors, and switches has proved well suited to infer high-level behaviours from low-level sensory data. Some systems such as the ones developed by Munguia et al (12) and Wilson & Atkenson (16) have used this approach obtaining good results using cheap sensors that can be deployed at low cost inside a home environment.

For our system, we have chosen the use of a temporal neural-network based embedded agent able to work with online, real-time data from unobtrusive low-level sensors and actuators.

The structure of this paper is as follows. In section 2, a brief introduction of the experimental testbed is presented, followed by the system's general architecture and the agent internal learning mechanisms. Section 3 addresses the experiments conducted and describes the results obtained from the system. Finally, the last section provides a brief summary of the work and outlines our future lines of research.

2. THE ACTIVITY RECOGNITION SYSTEM

2.1 Experimental Environments

Two different experimental environments have been built at the University of Essex. The first one is called the intelligent dormitory (iDorm1 – shown in Figure 1). This is a real pervasive computing testbed comprising a large number of embedded sensors, actuators, processors and networks in the form of a small self-contained room containing areas for different activities such as sleeping, communicating (writing or video conferencing with remote family and friends) and entertaining (watching TV, listening to music etc).

These networked devices enable the intelligent agents to monitor and make changes to the room's environmental conditions. The sensor network includes devices such as: temperature sensors (both inside and outside the room); humidity sensors; a small matrix of light sensors across the room; an active entrance lock system which provides access based on an individual's identity; an infrared sensor to detect movement. Effectors include: air circulators; fan heaters; a door lock actuator; motorised vertical blinds; automated window openers and a light dimmer.

FIGURE 1 - The iDorm1

With the success of the iDorm1, a more realistic test-bed for exploring care applications of ambient intelligence in the home has been constructed. The new facility takes the form of a domestic apartment and has been called the iDorm2.

The new intelligent dormitory 2 (iDorm2) shown in figure 2, is a full-size two bedroom apartment. This apartment is built from the ground up to support experimental work and features specially constructed cavity walls and ceilings that house sensors, effectors, processors, networks and power systems, all hidden from view and fully configurable by the researchers. All the basic services are electrically controlled (e.g. heating, water, doors etc).

Thus, the iDorm2 offers the possibility for examining the deployment of embedded agents and sophisticated user interfaces within the intelligent environments of tomorrow.

FIGURE 2 – The iDorm2

2.2 Agent description

The iDorm devices are interconnected to the agent using the Universal Plug&Play (UPnP) architecture which provides a component based API, service discovery and the communication of actions (events) inside the iDorm. The system is programmed using the Java language which brings advantages such as flexibility, modularity, reusability and portability.

The agent is able to run on a SNAP board which is a network-ready, Java-powered plug & play platform developed by Imsys (18).

The main purpose of running an agent on a SNAP board was to show that the agent could run on a real embedded-internet device which has, typically, an order of magnitude less memory and processing speed.

One of the major concerns related to care environments is the lack of proper agent architectures (both the internal & external agent structures and mechanisms) able to cope with the challenges of more demanding scenarios. Our work addresses this issue by exploring the use of a neural-network based agent in order to detect, recognize and classify human activities and behaviours inside an environment   In order to use a neural network for this purpose, the network must be able to identify recurrent patterns of behaviour, yet flexible enough to adapt itself to continuous changes in the environment.

Our approach comprises the use of an Adaptive Neural Architecture derived from the ECoS paradigm proposed by Kasabov (7).   This kind of network can grow dynamically, adapting its hidden layer to accommodate new information by adding nodes (rule nodes) whenever an example is not found to fit the existing structure. It is even able to "grow" new input or output nodes to

accommodate, for example, new sensor inputs or new activities.

Many of the abnormalities in behaviours that can be detected are not only related to the appearance of new activities but also to the temporal order in which they take place. With the addition of memory structures, the learnt temporal associations can be used to support the activation of the rule nodes based on temporal patterns, along with the existing spatial-similarity associations found in activities and human behaviours.

A complete description of the construction algorithm used by the network to update its internal weights can be found in previous work by Rivera-Illingworth et al (15). The following paragraphs will focus on the description of the memory layer and the changes made to the network in order to accommodate this layer.

2.2.1 TEMPORAL COMPONENTS OF THE NETWORK.

The vast majority of the artificial neural networks only deal with problems whose nature is "static", however time is a key component of human behaviours. Usually, a neural network is presented with an input pattern and after some processing (using feed-forward networks, perceptrons, etc.) an output is generated. These networks associate each one of the input patterns to a single output pattern.

Conventional neural architectures are not well suited time varient patterns, that is, for temporal pattern recognition that involves processing of patterns that evolve over time. According to Mozer (10), for this kind of pattern recognition, the appropriate response at a particular point in time depends not only on the current input, but potentially all previous inputs.

A popular way to recognize patterns that vary across time has been to use a recurrent network. Chappelier et al (3), regard recurrent neural networks (RNN) as a major family of temporal connectionist models. A RNN can be defined in the most general manner as a neural network containing at least one neuron whose state depends either directly or indirectly on at least one of its anterior states. According to Kremer (8), this characteristic provides a short-term memory which allows these networks to deal with input and output patterns that vary across time .

In a RNN, the connections are mainly feed-forward, but include a carefully chosen set of feedback connections. The recurrence allows the network to remember cues from the recent past, but does not complicate the training. If the feedback connections are fixed and not trainable, back-propagation may be easily used for training. The updating is synchronous, with one update for all units at each time step (5).

The RNN architecture proposed by Elman (4), is a recurrent network based on the MLP architecture. It consists of adding recurrent links from the hidden layer of the network to the input layer. At the time $t - 1$ the hidden layer is copied into the input layer as a complement to the "real" input vector at time t. When computing the new output, the information goes downstream, as in a classical MLP, from the input to the output layer (3).

More specifically, the network uses the previous state (also called context) together with the current sensory input as the input to the neural network and produces a new state as output. The short term memory is stored in a set of hidden units whose activations are computed based on the activations in a layer of input units, a layer of context units and on the weights from these to layers to the state units.

By convention, the activations of all the context units are initially set to 0.5 and subsequently set to the activation values of the hidden units at the previous time step (the number of context and hidden units must be equal). Specifically, the hidden unit activation values are copied to the context units at each time step. Thus, at any given time, the hidden units' activations represent the current state, whilst the context unit activations represent the previous state (8).

In this kind of architecture, the memory uses states based on the previous state vector and input vector. This implies that the state vector for this type of memory can contain information not found in recent input and output vectors.

An adaptation of an Elman (a Recurrent Neural Network) architecture was chosen because recurrence allows the network remember information from the recent past and does not appreciably complicate the training. The addition of a temporal layer and the connection weights allows the network to capture temporal dependencies between consecutive data examples.

Figure 3 provides a diagram of the actual structure of the adaptive neural network, consisting of four layers: the input layer, the evolving hidden layer and the output layer, plus a memory layer used to represent temporal information. The dashed arrows represent fully connected layers whereas the solid line arrow represent one-to-one connection between neurons.

FIGURE 3 - The network architecture
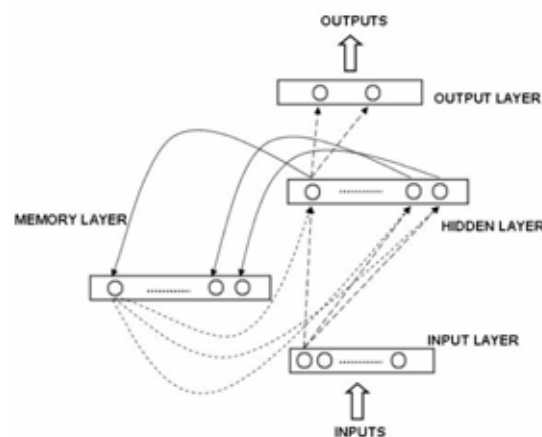
According to Elman (4), the memory layer is a structure that allows the network to have temporal capabilities. The network has feedback connections from the hidden

layer of neurons back to the same nodes, and these context units simply hold a copy of the activations of the hidden units from the previous time step. These context units or memory layer allow the network to capture temporal dependencies between the presented data examples from the data stream.

The activation function A of each neuron in the hidden layer is now dependant of both the spatial and temporal components. The proportion in which these two components influence the neuron activation can be modified by the spatial and temporal factors (*Sf* and *Tf* respectively). The activation function is calculated as shown in Equation (1):

$$A_{hnode} = 1 - \left(Sf * D_{I,hnode} + Tf * Wmh_{\max actv, hnode} \right) \quad (1)$$

In this $D_{I,hnode}$ represents the distance function between the input and a hidden node whereas $Wmh_{maxactv,hnode}$ represents the memory-to-hidden layer connection weight between the maximum activation neuron and the hidden node.

The computation of the input-to-hidden layer connection weights *Wih* and the hidden-to-output layer connection weights *Who* remains the same. The new memory-to-hidden layer connection weights *Wmh* capture the temporal dependencies between consecutive data examples.

These *Wmh* weights are updated using Equation (2), where A_h is the activation of the winning hidden node, A_m is the activation of the winning memory node and η_3 is the learning rate.

$$Wmh(t+1) = Wmh(t) + \eta_3 \left(A_h \cdot A_m \right) \quad (2)$$

Many of the abnormalities that can be detected are not just related to the appearance of new activities but also to the temporal order in which they take place. With the addition of the new memory structures, the learned temporal associations can be used to support the activation of the rule nodes based on temporal patterns together with the existing spatial-similarity associations found in the activities and human behaviours.

In the next section the experiments conducted in order to test both the ability of the network to find normal patterns of behaviour, and the possibility to detect deviations in the temporal order of those patters, will be presented.

3. EXPERIMENTS AND RESULTS

This section explains the type of experiments conducted to test the behaviour recognition system. The aim was not just to identify patterns in the sensors' activation, but also to detect high-level behaviours such as eating, using a computer, etc. The data was collected using the system described in the previous sections. The results presented show both the ability of the system to detect

and recognize human behaviours and also the possibility to detect deviations from those habitual patterns of behaviour.

3.1 Data acquisition

The dataset consisted of the data collected by 18 sensors (both environmental and furniture based). A set of eight different activities (Listening to music, Working at Computer, Reading, Desk work, Resting-napping, Sleeping, Out of room, and 'Other') were detected inside the environment. To assist this process, during the experiments, the user was asked to describe the action he was performing via a simple user interface, using an approach similar to the Experience Sampling Method (ESM) used in Munguia et al (12).

In order for the agent to collect the data related to the user's behaviour and activities inside the environment, two different approaches were used. In the first approach the agent recorded a 'snapshot' of the current inputs (sensor states) only when the user implicitly changed the state of the environment (e.g. turn on/off a switch); however, by using this approach, the number of instances recorded was limited (about 100 a day). By using this approach one dataset of 471 instances was collected over a period of 4 days.

For the second approach, a new module was added to the agent so the state of the environment (and the current action being done by the user) could be recorded every 30 seconds regardless of the user interaction with the system. In this way, more detailed information could be obtained; a second dataset consisting of a total of 1066 instances was recorded using this approach over a period of 2 days.

3.2 Experiments

The experiments were divided into two parts. The aim of the first part was to determine the performance of the new temporal network on classifying the activities and behaviours. The second part corresponded to experiments conducted in order to test the new temporal abilities of the network to identify time-based abnormalities.

3.2.1 CLASSIFICATION OF ACTIVITIES USING TEMPORAL NETWORK ARCHITECTURE. These series of experiments were conducted in order to compare the performance of the normal network and the one with temporal capabilities. For this experiment, both the first and second dataset were used, merging the 4 days of data of first dataset with the first day of data of the second dataset and testing over the second day of data of the second dataset. This was done also to test the performance of the network using data acquired in two different periods of time. The results for this test is summarized in the table below.

TABLE 1 – Comparison of classification between the normal and temporal network

| Training		Not temporal	Temporal
	Training examples	996	996
	Correctly classified	926	915
	Percentage	92.97%	91.86%
	Created hidden nodes	178	158
Testing	Test examples	541	541
	Correctly classified	492	500
	Percentage	90.94%	92.42%

As it can be seen from table 1, the temporal network has a performance as good as the normal network, and a better one for the testing phase. Although the temporal network created less hidden nodes, it took longer to train because the addition of the temporal layer meant that it had to train more inter-layer connections. The results also show that combining old and new data (acquired at different periods of time) of the same individual did not affect the classification results, so the network is able to generalize its classification results over different periods of time.

3.2.2 TEMPORAL ABNORMALITY DETECTION. Two different experiments were performed in order to test the ability of the network to identify abnormalities related with time. In the first experiment, the network was tested to detect an activity that it has been previously trained but, that this time, presented at a different hour of the day. A graphical representation of the normal order of the activities (Figure 4) and one with the activity taken place at a different hour (Figure 5) are shown in the figures below (with the hour in the Y-axis an the activities on the X-axis).

FIGURE 4 – Normal pattern of activities
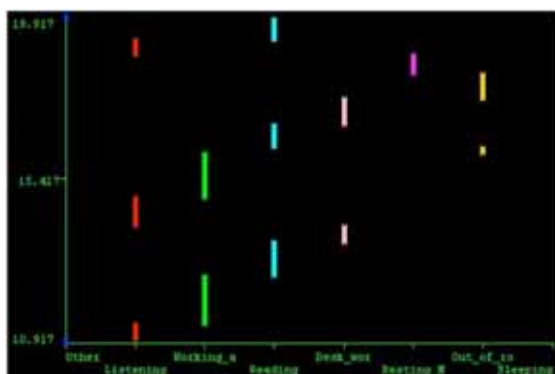
FIGURE 5 – Activities taking place at a different time
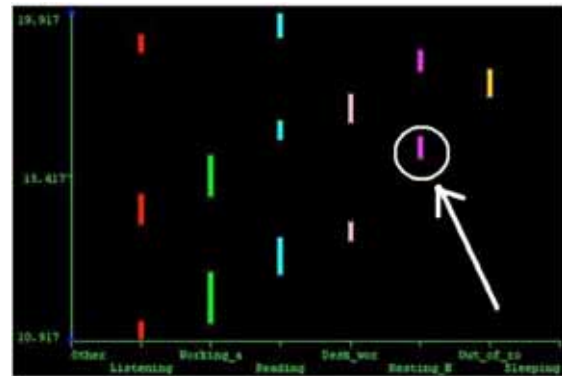
The network was tested in order to see if it can spot an activity taking place at an unusual time according to the normal pattern of behaviour. This abnormality consisted of the inhabitant taking a nap at a different hour. The results are summarized in table 2.

TABLE 2 – Abnormality detection for activities occurring at different time

| Testing	Novelty sensitivity thr.	0.95
	Detected novelties	31
	True novelties	31
	Mislabelled novelties	0

The network found 31 abnormal activities corresponding to the 31 instances in which the napping activity occurred at a different time, correctly spotting 100% of the abnormal instances. To get more conclusive results, further experiments need to be conducted to test more activities that have been previously seen but performed at different times of the day. However, from our initial results, the network seems to be well suited for this kind of task.

For the second experiment the network was tested to identify abnormalities relating to activities occurring in a different order. Again, the comparison again was made using a normal set of activities shown in Figure 4 and the new order of the activities can be seen better in the graphical representation presented in Figure 6.

TABLE 3 – Abnormality detection for activities occurring at different order

Testing	Novelty sensitivity thr.	0.98
	Total novelties	182
	Detected novelties	184
	True novelties	122
	Mislabelled novelties	62

FIGURE 6 – Activities taking place with different order
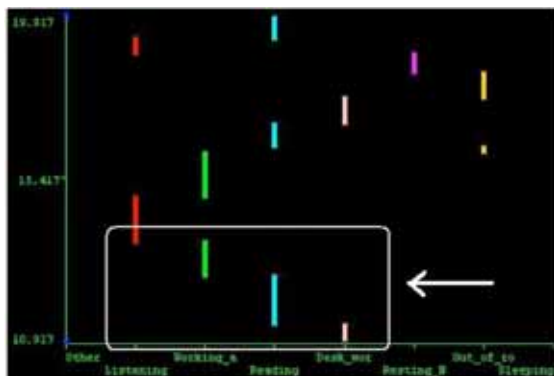
As it can be seen in table 3, the network found 122 true abnormalities, corresponding to 67% of the total real abnormalities. However, almost all the abnormalities that were not found corresponded to the "Working at computer" class. We believe that the reason for that was that this activity is sometimes difficult to distinguish due to the sensory data being similar to that of other activities.

The results of the two experiments have shown that the temporal network is able to find abnormalities in which the time plays a significant role. New sets of data will be collected so the network can be fully tested.

4. CONCLUSIONS AND FUTURE WORK

Many of the abnormalities that can be detected are not only related to the appearance of new activities but also to the temporal order in which those take place. With the addition of the new memory structures, the learned temporal associations can be used to support the activation of the rule nodes based on temporal patterns together with the existing spatial-similarity associations found in activities and human behaviours.

The recurrent neural network architecture proved to be a simple yet useful approach because it could be incorporated to the existing network without major modifications in the training algorithm.

The system is able to learn in an online (one-pass) and incremental way, adapting itself to new data as they are made available over time and it can be applied to environments with changing dynamics. Our system is able to perform as well as more complex systems, needing only to be trained for one epoch (the examples need to be presented only once to the network in order to train it) and uses an abnormality detection method that is embedded into the structure of the network which proves that the network can be expanded as new examples are presented or it can add new classes to accommodate the abnormal instances.

To our knowledge, there aren't any systems able to recognize human behaviours that can be integrated wholly within a computationally lean embedded processor. Thus, this is an important contribution of this work.

The system presented in this paper, presents an original approach to the activity detection and recognition of patterns in behaviour using data from unobtrusive and relatively simple sensors and the use of an adaptive neural network enhanced with temporal capabilities. This system is not only able to recognize high level behavioural activities such as "eating", "sleeping", "listening to music" but also it can recognize patterns on those behaviours. By using those patterns it can identify changes in the temporal order of the human activities (both in the order and in time of occurrence) that later on can be used for care applications.

The results of the use of the network with a temporal layer are encouraging and have shown that the abnormality recognition process benefits from the use of the temporal components by allowing the network to keep an internal memory. The recognition of patterns in behaviour also benefited from the use of more environmental data. Further experiments will be conducted in order to extend our results and test the system in bigger environments.

REFERENCES

1.   Abowd, G., Mynatt, E., Todden, T.: The Human Experience Reaching for Weiser's Vision, IEEE Pervasive Computing, 1(1):48-57, January-March 2002

2.   Beaudin, J., Intille, S., and Munguia, E., Lessons learned using ubiquitous sensors for data collection in real homes CHI Extended Abstracts 2004 1359-1362

3.   Chappelier, J., Gori, M., and Grumbach, A. : "Time in Connectionist Models", Lecture Notes in Computer Science, vol. 1828, pp. 105-?, 2001.

4.   Elman, J. L. : "Finding structure in time", Cognitive Science, vol. 14(2), pp. 179-211, 1990.

5.   Hertz, J., Krogh, A., and Palmer, R.G. Introduction to the Theory of Neural Computation, Addison-Wesley Publishing Company, California (1991)

6.   Holly B. Jimison, Misha Pavel, Jesse Pavel, "Adaptive Interfaces for Home Health", UbiHealth 2003: The 2nd International Workshop on Ubiquitous Computing for Pervasive Healthcare Applications, Seattle, WA, Oct 12, 2003

7.   Kasabov, N. Evolving connectionist systems: Methods and applications in bioinformatics, brain study and intelligent machines, Springer,London (2002)

8.   Kremer, S. C. : "Spatio-temporal Connectionist Networks: A Taxonomy and Review", Neural Computation, 13(2):249-306, February 2001.

9.   Mihailidis, A., Carmichael, B., Boger, J., and Fernie, G.: "An Intelligent Environment to Support Aging-in-Place, Safety, and Independence of Older Adults with Dementia", Ubi-Health 2003. Seattle, WA, Oct 12, 2003.

10.   Mozer, M. C. : "Neural net architectures for temporal sequences processing" In, Time Series Prediction, p.p. 243-264. Addison--Wesley, 1994.

11.   Mozer, M.: "The Neural Network House: An Environment that Adapts to its Inhabitants", In M. Coen (Ed.), Proceedings of the AAAI Spring Symposium on Intelligent Environments, pages 110--114, 1998

12.   Munguia, E., Intille, S., and Larson, K.: Activity Recognition in the Home Using Simple and Ubiquitous Sensors. Pervasive Computing, Second International Conference, PERVASIVE 2004, Vienna, Austria, April 21-23, 2004: 158-175

13.   Mynatt E, and Rogers, W "Developing technology to support the functional independence of older adults" Ageing International, 2002

14.   Philipose, M., Fishkin, K., Perkowitz, M., Patterson, D., Fox, D., Kautz, H., and Hähnel, D.: " Inferring Activities from Interactions with Objects". IEEE Pervasive Computing, October 2004, pp. 50-57

15.   Rivera-Illingworth, F.; Callaghan, V.; Hagras, H.: A connectionist embedded agent approach for abnormal behaviour detection in intelligent health care environments. Systems, Man and Cybernetics, 2004 IEEE International Conference on Volume 4, 10-13 Oct. 2004 Page(s):3565 - 3570 vol.4

16.   Wilson, D., and Atkeson, C.: Simultaneous Tracking and Activity Recognition (STAR) Using Many Anonymous, Binary Sensors. In Proceedings of PERVASIVE 2005, Munich, Germany, May 2005

17.   Zita Haigh, K., and Yanco, H.: "Automation as Caregiver: A Survey of Issues and Technologies", AAAI-02 Workshop on Automation as Caregiver: The Role of Intelligent Technology in Elder Care, pp. 39-53. July 2002

18.   http://www.imsys.se/