Unsupervised Learning in Ambient Assisted Living for Pattern and Anomaly Detection: A Survey

Francisco Javier Parada Otte, Bruno Rosales Saurer, and Wilhelm Stork

FZI Research Center for Information Technologies
Embedded Systems and Sensors Engineering Karlsruhe, Germany
parada@fzi.de

Abstract. Population ageing is an issue that has encouraged the development of Ambient Intelligence systems to support elderly people to live autonomously at home longer. Some key aspects of these systems are the detection of behavior patterns and behavior profiles in their daily life. The information we can infer from these patterns could prove to be very valuable for monitoring the health status of a person, like to control deterioration of diseases or to provide personalized assistive services. In this paper we focus on the unsupervised learning techniques in health monitoring systems for elderly people, which has the advantage of not needing annotations. Collecting these is a tedious job and sometimes difficult to accomplish. We discuss the different existing approaches, identify some limitations and propose possible challenges and directions for future research.

Keywords: Ambient Assisted Living, unsupervised learning, behavior profiling, health monitoring.

1 Introduction

Ambient Assisted Living (AAL) is the application of Ambient Intelligence (AmI) to support elderly people to live at home longer. Population ageing has become a very important fact in our society due to the social and economic implications that it entails [1].

An AmI system can be dissected in three main parts [2], as seen in Fig. 1: sensing, reasoning and acting. *Sensing* is the ability to capture or record any type of context information that is relevant for the system. This could be any information about the user and his environment. This information is collected mostly through sensors, which provide a wide variety of data, but there could also be user interfaces to get user input or some predefined information. Sensors could be environmental like cameras, microphones, motion, contact, brightness, and temperature sensors, or wearable like accelerometers and vital parameters sensors.

The second part, *reasoning*, consists in obtaining more abstract or higher level information from the sensor data in order to acquire valuable and useful information for the system. The process is based on the application of some inferring techniques to the context data to obtain the information. There are mainly two most extended types;

M.J. O'Grady et al. (Eds.): AmI 2013 Workshops, CCIS 413, pp. 44–53, 2013.

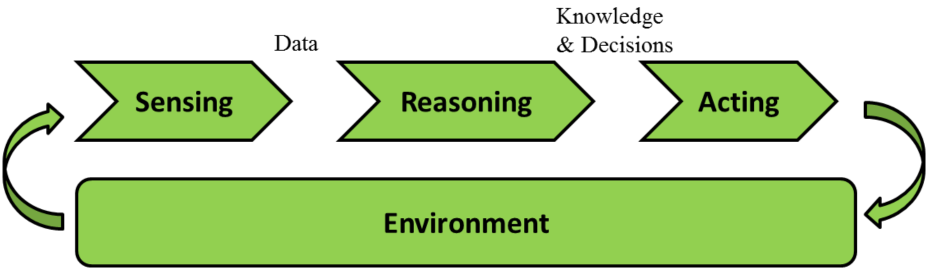

Fig. 1. Structure of AmI Systems [2]

on the one hand, we have specification-based systems, where some domain knowledge is reflected in defined rules for inferring the information. On the other hand, we have machine learning, which explores the relationships between the context data and the information to be inferred. Machine learning can be divided as well into two subcategories, supervised and unsupervised. The difference lies in the usage or not of labeled data for the learning process respectively. Supervised learning needs annotations or references in order to train the system to learn the relationship model. But the process of collecting the annotations is very tedious and time consuming. Sometimes it is very difficult to carry out such tasks, for example, in large installations, in long-term studies and with people who have cognitive problems. In these cases unsupervised learning can be applied.

The third part, *acting*, basically defines the application or goal of the system. Here the system reacts to the information inferred in a proper way. It could present the information to the caregiver in case of health monitoring systems, provide a service if it is an assistive system, or inform the responsible person in case of emergency detections.

The analysis and detection of behavior patterns belong to the reasoning part and is a cornerstone in AmI systems, and especially relevant in AAL. The behavior patterns, how they evolve, and the detection of deviations from expected behavior can provide very valuable information about a person's health status. The main applications in the home environment are the health status monitoring over time, for detecting depression, control early stages of dementia or chronic diseases like diabetes. Secondly, it is possible to monitor the progress of patients after surgery, in rehabilitation or under treatments like for multiple sclerosis. Then, it could also work as a supporting tool for caregivers as backup for assessments like the Barthel index [3] and other documentation. The system is also useful for emergency detections like an inactivity period after a fall. And the last application to remark is the provision of personalized services in assistive systems, such as lighting control at night.

In this paper we examine the work performed in the field of monitoring systems for behavior patterns, activity recognition and anomaly detection, centering on the unsupervised learning techniques for AAL. In the next section we cover the other main alternatives to the unsupervised learning, the ones just mentioned above. We explain and discuss them compared to unsupervised learning. In the third section we discuss the current work based on unsupervised learning. We classify and comment the differences among them. Finally in the last section we discuss the current status of the research in this field and what might be interesting challenges and future directions for research for the community.

2 Specification-Based and Supervised Learning Techniques

Besides unsupervised learning, there are other reasoning techniques as previously mentioned, which we present now compared to unsupervised learning. It is important to acknowledge, that depending on what the goal of the system is, or what information is needed, different techniques can fit the requirements better.

Specification-based systems apply expertise knowledge in a domain in form of logic rules. Reason engines then infer the situations from the sensor data. Bikakis et al. provide an extended survey describing the work carried out in this direction [4]. They distinguish among three main types and some minor solutions. The first one is ontological reasoning, where an ontology represents the data model and the systems benefit from the tools already available like query languages and reasoning tools. Secondly, they mention the rule-based reasoning, where logic rules define policies, constraints and preferences applied to the context data in order to retrieve higher level information. Then the third model is distributed reasoning, where different reasoning components interact together. These components may have different reasoning, storage and computing capabilities, and may deal with data from different sources and different formats. These systems perform well when using small datasets, but they have limited computational capacity. The rapid development of pervasive devices increases the use of a more diversity of sensors, thus, the datasets increase as well, becoming more incomplete and uncertain. The efficiency and the capacity of retrieving information are considerably reduced in this case. A second drawback is the lack of adaptability of such systems. Different people will generate different behavior patterns and will also change over time, while these systems cannot learn from the changes. This technique is limited as a stand-alone system, but combined with learning algorithms, it could provide very interesting results.

Supervised learning deals better with the huge datasets and with the incompleteness and uncertainty. These techniques explore associations between the sensor data and the situations to be inferred. There is a lot of work in this field and Aztiria, Izaguirre and Augusto have surveyed these techniques centered on pattern learning in Ambient Intelligence [2]. Neural networks, classification techniques like decision trees or Bayesian networks, and sequence discovery are examples of the most extended algorithms. The main drawback of these systems is the need of labeled data for a training phase. The collection of annotations is a very laborious and time-consuming task and sometimes impossible to achieve. It also might not be reliable when elderly people annotate their own activities as they may become forgetful or have difficulties. There have been some works improving the collection, gathering them electronically via a PDA [5] or taken them from a recorded video [6]. However, doing this for each new installation would be arduous for large-scale evaluations or when bringing such systems to the market. For the matter of research of new algorithms, it is helpful to do a training phase with labeled data, but when considering real-world conditions, it would become hard to collect the annotations. Unsupervised learning could fit in these cases.

There exist two more approaches lying in between supervised and unsupervised learning. First, semi-supervised learning uses both labeled and unlabeled data for the learning process, reducing the amount of labeling data needed. They use the labeled data to label automatically the unlabeled data with some classifier, incrementing the labeled set for the model. There is some work in behavior modeling [7] and detecting unusual events [8]. This system is proper when only a few annotations are available, but these are still needed. The second alternative is reinforcement learning, where the learning process is supported by a feedback function, indicating how well the system operates. Thus, the model adapts itself according to the feedback received. Reinforcement learning works well for adaptive systems, like light control systems, using the user actions (switching on/off) as feedback [9]. In our case, focusing on finding outliers, we cannot rely on the feedback from the sensors, but the concept is interesting to apply adaptability to the system.

3 Unsupervised Learning

There is an active work in pattern recognition and outlier detection and the work in unsupervised learning is not less. It is difficult to classify the different current solutions, as the algorithms differ from each other. Most of them are based on cluster techniques, which are the most extended unsupervised techniques, but there are some others as we will see. Other solutions are combinations of two algorithms, where clustering techniques are often also involved. We have categorized them in clustering-based solutions and mixed solutions.

3.1 Clustering Techniques

Clustering consists mainly in dividing a dataset in subsets by grouping them with some similarity criteria. There are two different clustering techniques, partitional and hierarchical. In the former, clusters are found at once while in the latter, clusters are found iteratively based on previous clusters. Some concrete algorithms are density models, where clusters are the densest regions, distribution models based on statistical distributions, or centroid models, where all clusters are mean vectors and the belonging relationship is some distance parameter to the center value. K-means is a well-known centroid algorithm based on the partition of a dataset into k clusters in which each observation belongs to the cluster with the nearest mean. The term nearest is based on a metric, as it could be the Euclidean distance. Although clustering is the most common, it is often combined with other techniques to achieve better results and is not used as a stand-alone.

Monekosso applies unsupervised learning to find out routines based on daily activities [10]. She uses clustering techniques, concretely k-means, k-medoids, agglomerative clustering and EM (expectation-maximization) clustering. Then she built a behavior model based on a hidden Markov model (HMM), defining behavior as a sequence of clusters. For the evaluation she collected data from one household for several alternate weeks, and then applied them to the algorithm. It is a good start, but validation of this work in real-time and with more subjects would be desirable, as these systems should work in real-world conditions. It is quite interesting, that she evaluates all four algorithms for pattern detection with HMM and different hidden states, and then for anomaly recognition. Results vary depending on the number of hidden states and the length of the sequence. This addresses the difficulty of choosing the best algorithm, as the performance varies depending on several factors. Behavior not recognized with the model is tagged for informing caregivers and they receive some information to make decisions. This information is the probability of an anomalous behavior, and the behavior itself, like a lack of activity in a room. This output would result in useful information for a monitoring system.

Barger and the group at the University of Virginia study how to obtain behavioral patterns for a monitoring system in a Smart House by applying mixture-models [11]. The so-called mixture models combine the k-means clustering technique and self-organizing maps, which is a sort of neural network, as a method for event estimation. For their study they use real motion data from a Smart Home collected discontinuously during 65 days, 25 for training and 40 for the tests. They make the distinction in their system between work and days off by activity level, thus, the amount of sensor firings, which helped to reduce the uncertainty when classifying the observations. It also allows for the possibility of finding different patterns in work and days off. The evaluation is then performed offline including the unsupervised training phase. A user log was available as well, allowing them to compare the patterns to the activities performed by the user. They could identify many of the recorded patterns as activities such as sleep behavior, changing clothes or bathroom use. The output of the system would be sequences of motion sensor firings, which are hard to interpret by the end user.

The MavHome group is a very active working group in pattern recognition and behavior prediction in intelligent environments [12]. Regarding clustering techniques, Rao and Cook use the task-based Markov model which consists of the combination of clustering and HMM [13]. The system aims to predict a person's behavior in an intelligent environment, in order to provide assistance and adaptation to the inhabitant's needs. They divide the sensor data in groups based on some heuristic rules, then apply clustering to these sequences to represent tasks, and finally the HMM is applied for the behavior prediction. This group validated their system offline with simulated and real data, which provided different results. This fact emphasizes the importance of using real data. One positive aspect is that they compared their algorithms with a simple Markov model, showing a better performance for certain datasets. They state that the choice of the algorithm parameters (number of clusters, sequence length and allowable time difference) is important for the algorithm performance and they vary for concrete datasets. This issue shows the algorithms are closely dependent on several factors such as the configuration parameters or the data format and completeness, implying the difficulty of finding a universal algorithm that best fits all situations. The output information is a set of patterns of sensor data, which are usable for the system itself, but difficult to interpret by non-technical experts.

Fuzzy logic is also used as an unsupervised technique in the smart environment iDorm [14]. Their end goal is to provide an automatic control of the intelligent environment based on the user's behavior. They propose an online adaptive framework for extracting fuzzy membership functions and rules representing the behavior of persons. These functions are extracted directly from the sensor data by applying a double clustering approach. The system adapts to change behaviors thanks to a loopback module. They performed an online and offline evaluation with data collected from the iDorm during 5 consecutive days. The offline experiment consists of the performance comparison among their solution and 3 more algorithms, namely GP (genetic programming), ANFIS (adaptive-neuro fuzzy inference system), and the MLP (multiplayer perceptron), a neural network. Although their solution is the second in results after GP, it provides a lighter computation than the other ones and it is better suited for online operation. The online evaluation showed during 2 days, after 3 of learning, that the system can learn the functions and respond appropriately with the control rules. Again, the output is based directly on sensor data, hard to interpret by non-technicians, as would be expected for a monitoring system.

3.2 Other Unsupervised Solutions

Besides clustering techniques, there are several different solutions that perform unsupervised learning that apply a variety of algorithms for discovering behavior patterns.

Wyatt and Intel Research people were pioneers in activity recognition based on unsupervised learning [15]. They focused on labeling and learning sensor data from information obtained by web mining processes without human intervention to recognize activities but not behavior patterns. They mined from the web activities description to get a correlation of appearance of objects usage. Their experiments were offline with real data from 9 subjects, each subject only performing activities for 20 to 40 minutes in an equipped home. Nevertheless, this solution works only when every object is tagged, which would result in a very complicated process, especially for long-term studies. It would also be interesting to test it with ambient sensors, although the mining process should be adapted to this sensor input. It is also dependent on how good, quantitative and qualitative, the available descriptions are, and there is no possible personalization. However, it is a very interesting work, as it provides a completely different perspective when avoiding human intervention. The output presented is already at an activity level instead of a sensor level, thus easier to understand by non-experts.

Robben has some recent work on pattern discovery for health status monitoring. This group presents the possibility of predicting functional health status from ambient sensor data [16]. They applied expert knowledge for extracting important features for the caregivers and then looked for mapping the features to the functional status by linear regression and Gaussian processes. The feature extraction is matched to the

AMPS (assessment of motor and process skill), a functional health metric [17], which was performed and collected by caregivers. They carried out an offline experiment with data from 9 subjects at an assisted living facility and different time lengths between one and five weeks. They have found that some features can be detected, but individual differences are very influential and should be taken into account. Results are quite interesting and show that at this early stage it is not possible to generalize and results vary among the participants. They also present a visualization model of deviations in ADLs [18], which is a key aspect in monitoring systems. This information shown here is still sensor based and can be abstracted to a more understandable format for caregivers.

Virone works on detecting long-term activity patterns based on the circadian activity rhythm (CAR) [19]. The study of temporal structures called chronobiology, has demonstrated that human behavior is ruled approximately by a 24-hour fluctuating rhythm, known as the circadian rhythm. He uses statistical software called SAMCAD for detecting behavior patterns. He estimates the average time spent in each room and the average number of motion events per room, named as activity, in 24-hour periods. Anomalies result from comparing the behavior of the current day and the habitual behavior estimated on the basis of the CAR. These might be indicative of a problematic situation or a change in the person's health status. Results are sensor data estimations which could lead to very helpful information. Nevertheless, it is still presented in graphics of activity, as seen in Fig. 2, difficult to understand by non-experts. The detection is done in an offline process, although he aims to evaluate it in real-time for further work.

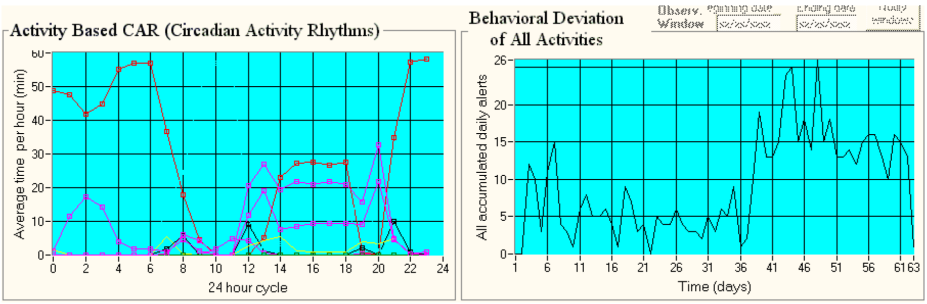

Fig. 2. Application for caregivers [16]

There are some more works using similar algorithms. A combination of HMM and a Viterbi algorithm is applied to analyze motion sensor data to model elderly persons' behavior and find unusual activities to inform caregivers [20]. Then, unsupervised neural networks have been applied as well, to find patterns in the Adaptive Home, a Smart Home environment, for improving comfort by home automation [21]. Rivera also explored adaptive neural networks for anomaly detection again in the iDorm intelligent environment [22].

4 Discussion and Research Challenges

There is a lot of work going on in pattern recognition and unsupervised learning is a very interesting technique, especially for anomaly detection, which might prove to be very useful for monitoring elderly people at home. No need for labeled data is also an advantage for large-scale evaluations or bringing a system to the market. However, there is still some work ahead to get full functional monitoring systems.

One of the limitations of unsupervised learning is the absence of the interpretation of the outputs. This is of special interest in monitoring systems, where readable information should be presented to the caregivers or the person responsible for the monitoring. The results and information that we get for such algorithms are not understandable for the end user. A pattern of sensor events or a deviation of this pattern is far from what a caregiver could expect. That is why many works combine two different techniques, so they can identify activities and apply the algorithm to an activity level and not only to the sensor level. One approach that could be interesting is to apply some semantic knowledge in order to interpret the results of the algorithms and to provide valuable information. Thus, when an outlier is detected due to a lot of motion at night time in contrast to the normal behavior, which is just light motion in the bedroom, the system could report some sleeping problems on a concrete date. However, there could be a large difference in the outliers and people also have different behaviors, which would complicate this approach. One option would be to implement algorithms that provide a readable output, or as said before, combine different techniques. We consider there is a lot of work in interpreting and visualizing the information, in order to provide a useful system for the end users.

On the other side, every work presents different algorithms and there is a wide variety of solutions. Each work uses their own data, real or simulated, and performs the tests and evaluation with it. This makes it very hard to make comparisons among the different systems. A good example to follow is the Artificial Intelligent community, where they make datasets available, allowing other techniques to be evaluated with these same datasets. There already are some public datasets, but still not really extended and used. Another proposal that would help to work in this direction is to standardize the sensor data format. This would facilitate sharing datasets and comparing solutions so everybody could benefit. Currently there are not any standards extended in the community.

Another aspect of the evaluation is the time when it is performed. Most of these evaluations are carried out offline. This means the data is collected for a specific time period or even using simulated data, and then the complete dataset is passed to the algorithm. In any case, the main purpose of these systems is to run independently and be able to work and respond in real-time. Simulated data is also good for a first evaluation, but this system should be validated both with real data and in real-time. We have seen in one work before, how different the results are when using simulated or real data [13]. Currently it is difficult to have the opportunity to perform such an evaluation, but infrastructures, like smart home environments, which are increasing quickly, and investment in this field, are making these studies become a reality.

Another issue when evaluating the systems is the duration of the datasets. So far only short time periods have been used, with a maximum of a few days or even a few hours. From this data, some is also used for training, leaving even less data for the evaluation. However, health monitoring systems aim to provide behavior patterns which should be recognized for long-term periods. Thus, some studies or datasets for longer time periods, for instance from one month in advance, would be desirable in order to validate such techniques. Secondly, online studies should be performed as well, as we addressed before. Although it is hard to make such studies and get subjects for a long-term period, there are already some institutions that allow carrying out these studies and there is an increasing interest for actors involved in eldercare like public administrations and healthcare organizations.

Finally, we also encourage the implementation of adaptive systems. A system should not only learn at the first stage and then keep working with the learnt model, but it should also keep learning over the time, as a person's behavior can change easily. This is already present in some works, but we think it is an essential feature for such systems.

There is still some work to be done and we propose here some directions for future research, but each contribution brings the community some advance in the field. This evolution will allow us to get improved AmI systems to support elderly people and everyone else involved in their care.

References

1. Bloom, D.E., Canning, D., Fink, G.: Implications of population ageing for economic growth. Oxford Review of Economic Policy 26(4), 583–612 (2010)
2. Aztiria, A., Izaguirre, A., Augusto, J.C.: Learning patterns in Ambient Intelligence environments: A Survey. Artificial Intelligence Review 34(1), 35–51 (2010)
3. Mahoney, F.I., Barthel, D.: Functional evaluation: the Barthel Index. Maryland State Medical Journal 14, 56–61 (1965)
4. Bikakis, A., Patkos, T., Antoniou, G., Plexousakis, D.: A Survey of Semantics-based Approaches for Context Reasoning in Ambient Intelligence. In: Mühlhäuser, M., Ferscha, A., Aitenbichler, E. (eds.) AmI 2007 Workshops. CCIS, vol. 11, pp. 14–23. Springer, Heidelberg (2008)
5. Stumpp, J., Anastasopoulou, P., Sghir, H., Hey, S.: Sensor Chest Strap Wirelessly Coupled with an e-Diary for Ambulatory Assessment of Psycho-Physiological Data. In: 2nd Biennial Conference of the Science of Ambulatory Assessment, Ann Arbor, Michigan (2011)
6. Tolstikov, A., Biswas, J., Tham, C.K., Yap, P.: Eating activity primitives detection - a step towards adl recognition. In: 10th IEEE International Conference on e–Health Networking, Applications and Service, HEALTHCOM 2008, pp. 35–41 (2008)
7. Hoey, J., Poupart, J., Boutilier, C., Mihailidis, A.: Semi-supervised learning of a POMDP model of Patient-Caregiver Interactions. In: International Joint Conference in Artificial Intelligence, Workshop on Modeling Others from Observations, pp. 101–110 (2005)
8. Zhang, D., Gatica-Perez, D., Bengio, S.: Semi-supervised adapted HMMs for unusual event detection. In: Conference Computer Vision and Pattern Recognition, CVPR 2005, pp. 611–618 (2005)

9. Sandhu, J.S., Agogino, A.M., Agogino, A.K.: Wireless Sensor Networks for Commercial Lighting Control: Decision Making with Multi-agent Systems. In: Association for Advancement of Artificial Intelligence Conference, Sensor Networks Workshop, San Jose, pp. 88–92 (2004)
10. Monekosso, D.N., Remagnino, P.: Anomalous behaviour detection: supporting independent living. In: Monekosso, D.N., Remagnino, P., Kuno, Y. (eds.) Ambient Intelligence Techniques and Applications, Advanced Information and Knowledge Processing, pp. 33–48. Springer, London (2009)
11. Barger, T.S., Brown, D.E., Alwan, M.: Health-status monitoring through analysis of behavioral patterns. IEEE Transactions on SMC-A 35, 22–27 (2005)
12. Cook, D.J., Youngblood, M., Heierman III, E.O., Gopalratnam, K., Rao, S., Litvin, A., Khawaja, F.: MavHome: An Agent-Based Smart Home. In: First IEEE International Conference on Pervasive Computing and Communications, PerCom 2003, pp. 521–524 (2003)
13. Rao, S., Cook, D.J.: Predicting Inhabitant Actions Using Action and Task Models with Application to Smart Homes. International Journals of Artificial Intel. Tools 13(1), 81–100 (2004)
14. Doctor, F., Hagras, H., Callaghan, V.: A fuzzy embedded agent-based approach for realizing ambient intelligence in intelligent inhabited environments. IEEE Transactions on Systems, Man, and Cybernetics, Part A 35(1), 55–65 (2005)
15. Wyatt, D., Philipose, M., Choudhury, T.: Unsupervised activity recognition using automatically mined common sense. In: 20th National Conference on Artificial Intelligence, Pittsburgh, Pennsylvania, pp. 21–27 (2005)
16. Robben, S., Krose, B.: Longitudinal Residential Ambient Monitoring: Correlating Sensor Data to Functional Health Status. In: 7th International Conference on Pervasive Computing Technologies for Healthcare, PervasiveHealth, pp. 244–247 (2013)
17. Fisher, A.G.: Assessment of Motor and Process Skills, 6th edn. Development, Standardization, and Administration Manual, vol. 1. Three Star Press Inc., Fort Collins (2003)
18. Robben, S., Boot, M., Kanis, M., Kröse, B.: Identifying and Visualizing Relevant Deviations in Longitudinal Sensor Patterns for Care Professionals. In: 7th International Conference on Pervasive Computing Technologies for Healthcare, PervasiveHealth, pp. 416–419 (2013)
19. Virone, G., Sixsmith, A.: Monitoring activity patterns and trends of older adults. In: 30th IEEE Engineering in Medicine and Biology Society, Microtechnologies in Medicine & Biology, pp. 2071–2074 (2008)
20. Yin, G., Bruckner, D.: Daily activity model for ambient assisted living. In: Camarinha-Matos, L.M. (ed.) DoCEIS 2011. IFIP AICT, vol. 349, pp. 197–204. Springer, Heidelberg (2011)
21. Mozer, M.C.: Lessons from an Adaptive Home. In: Cook, D.J., Das, S.K. (eds.) Smart Environments: Technologies, Protocols, and Applications. John Wiley & Sons, Inc., Hoboken (2005)
22. Rivera-Illingworth, F., Callaghan, V., Hagras, H.A.: Neural Network Agent Based Approach to Activity Detection in AmI Environments. In: IEE International Workshop on Intelligent Environments, pp. 92–99 (2005)