Artif Intell Rev (2010) 34:35–51
DOI 10.1007/s10462-010-9160-3

Learning patterns in ambient intelligence environments: a survey

Asier Aztiria · Alberto Izaguirre · Juan Carlos Augusto

Published online: 23 May 2010

Abstract It is essential for environments that aim at helping people in their daily life that they have some sort of Ambient Intelligence. Learning the preferences and habits of users then becomes an important step in allowing a system to provide such personalized services. Thus far, the exploration of these issues by the scientific community has not been extensive, but interest in the area is growing. Ambient Intelligence environments have special characteristics that have to be taken into account during the learning process. We identify these characteristics and use them to highlight the strengths and weaknesses of developments so far, providing direction to encourage further development in this specific area of Ambient Intelligence.

Keywords Ambient intelligence · Intelligent environments · Pattern learning · Machine learning techniques

1 Introduction

Ubiquitous Computing, a term introduced by Mark Weiser (1991), refers to a paradigm where a new type of relation between users and technology is established such that technology is widespread and transparent to the users. An important follow up to this concept has been developed under the term Ambient Intelligence (AmI) (Ducatel et al. 2001; Aarts 2004; Shadbolt 2003).

A. Aztiria (✉) · A. Izaguirre
University of Mondragon, Mondragon, Spain
e-mail: aaztiria@eps.mondragon.edu

A. Izaguirre
e-mail: aizagirre@eps.mondragon.edu

J. C. Augusto
University of Ulster, Jordanstown, UK
e-mail: jc.augusto@ulster.ac.uk

AmI is sometimes confused with similar but not equivalent terms: Smart Environments (Cook and Das 2005) and Pervasive Computing (Friedemann and Mahmoud 2002). AmI gives more importance to the user, the idea of human-centered design and the intelligence needed to allow the system to anticipate needs of the user. This shapes the decision-making of these systems in a different way and can be summarized as: “a digital environment that proactively, but sensibly, supports people in their daily lives” (Augusto 2007). Equally, AmI should not be confused with Smart Homes (Augusto and Nugent 2006). Smart Homes are only one form of the many possible realizations of AmI, which include Smart Classrooms, Smart Offices, Smart Cars and others (see Augusto 2007; Augusto and Cook 2007a for a more detailed discussion). However, Smart Homes are so far the most extensively explored of those applications.

Let us consider the next scenario, which exemplifies typical activity development in an AmI environment: “Michael is a 65-year-old man who lives alone and enjoys an AmI system that makes his daily life easier. Usually, Michael’s alarm goes off around 07:50. Approximately 10 min after getting up, he usually steps into the bathroom; at that moment, the lights are turned on automatically and a screen installed in the bathroom reminds him what tablets he has to take. On weekdays, the shower is always ready according to Michael’s preferences (around 24–26°C in the winter). Before he leaves the bathroom, the environmental temperature has already been increased to 19°C, and all blinds are opened. He can watch the latest headlines about his favorite soccer team in the kitchen on a screen while he eats his breakfast. When Michael leaves the house, all the lights are turned off. Safety checks (e.g., checking to make sure the cooker is turned off) are performed to prevent hazardous situations in his absence, and, if needed, the house acts accordingly (e.g. by turning the cooker off)”.

Ambient Intelligence as a technological paradigm has the potential to make a significant impact upon daily human life by positively altering the relationship between humans and technology. In an AmI environment, “people would be surrounded by intelligent intuitive interfaces that are embedded in all kinds of objects and an environment that is capable of recognizing and responding to the presence of different individuals in a seamless, unobtrusive and often invisible way” (Ducatel et al. 2001).

The area has attracted a significant number of researchers, and some applications are already being deployed with different degrees of success. Taking into account the complexity of Ambient Intelligence systems (e.g., hardware, software and networks have to cooperate in an efficient and effective way to provide a suitable result to the user), each project has focused upon different aspects of such complex architectures. In that sense, it is understandable—and even logical—that the first developments have been focused upon needs associated with hardware and networking as supporting infrastructure. This has resulted in simple automation that implements a reactive environment. Giving more importance to the intelligence component is necessary, however, given its relevance for achieving the core aspects of AmI systems. Although many researchers to date have noted the importance of the topic (Augusto and Nugent 2006; Friedwald et al. 2005; Das and Cook 2006; Ramos et al. 2008) little emphasis has been placed in general upon the subject. However, notable exceptions can be found (see Sect. 4).

One of the hidden assumptions in AmI is that unlike current computing systems where the user has to learn how to use the technology, the environment in AmI environments adapts its behavior to the user. Ideally, the user should be released from the burden of programming any device (Hagras et al. 2004). The environment should learn how to react to the actions and needs of the user, and this should be achieved in an unobtrusive and transparent way. Thus, the ability to learn patterns of behavior becomes an essential aspect for the successful implementation of AmI systems.

In an AmI environment, learning means that the environment has to gain knowledge about the preferences, needs and habits of the user in order to better assist the user (Galushka et al. 2006; Leake et al. 2006).

Furthermore, it is insufficient to learn user patterns only once because preferences and routines can change with time. Hence, it is necessary for an AmI environment to adapt itself to these new patterns continuously (Rutishauser et al. 2005). Learning and adapting to user patterns is an essential feature of an AmI system. In a smart home used for healthcare, for example, it allows the system

- To understand common behaviors of the user so that relatives or caregivers can supervise and assess his/her patterns of behavior and detect unhealthy habits (e.g., by knowing how often and when Michael usually eats, the system allows his relatives to control his diet).
- To automate activation/deactivation of some devices depending on the needs of the user (e.g., by turning the lights on when Michael goes into the bathroom) (Cook et al. 2003; Heierman and Cook 2002).
- To make a system more efficient (Mozer et al. 1995) (e.g., by switching the lights off when Michael has gone out and will not return soon).
- To increase safety (Rivera-Illingworth et al. 2005; Liao et al. 2004; Kautz et al. 2002) (e.g., by either switching the cooker off or issuing an alarm when detecting that Michael has left it on).

As with any new proposed methodology, the AmI vision is not without criticisms. For example, there are concerns with regard to loss of privacy or fear of an increasingly individualized society. These criticisms have previously been applied to Computer Science as a whole. We are aware of the importance and complexity of users in AmI environments. Human aspects are very important and thus must be taken into account during the learning process.

On one hand, hundreds of papers and even books (Mitchell 1997; Witten and Frank 2005) relating to Machine Learning (ML) topics have been published. On the other, while some surveys on different AmI topics have been published (Cook and Das 2007; Pantic et al. 2006; Jiang et al. 2004), a survey relating learning skills to AmI environments is not available. This article aims to fill the gap in the use of ML techniques in order to learn common behaviors of the user in different environments. The remainder of the paper is organized as follows. Section 2 defines the necessary technologies to achieve real AmI environments. Section 3 describes the special features of AmI environments to be considered when performing the learning process. Section 4 reviews different approaches employed by different groups highlighting the strengths and weaknesses of the ML techniques they use. Finally, Sect. 5 provides some reflections on this topic.

2 Enabling technologies

The IST Advisory Group lists the key technologies that are required to make AmI environments a reality (Friedwald and Costa 2003). Among all the technologies they list, the learning process will be highly influenced by those technologies oriented to provide the environment with intelligence. These are sensing, reasoning and acting technologies. Sensing systems allow to perceive the state of the user and the environment by means of sensors, then, reasoning systems use that data to decide how to act upon the environment to get the intended goals, and finally, acting systems carry out these decisions. The process is illustrated in Fig. 1.

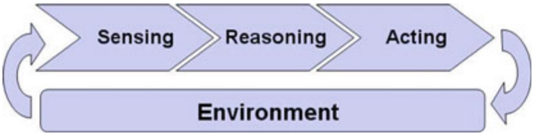

Fig. 1 Process of AmI systems

2.1 Sensing

The first step in the process of providing the environment with intelligence is to know the state of the user or users being supported by the system as well as the state of the environment itself. Sensors are a key technology that allows to link the real world the reasoning algorithms. There are many different types of sensors, for example for location tracking (Wolffenbuttel et al. 1990), for humidity sensing (Delapierre et al. 1983), for physiological sensing (Ermes et al. 2008; Stanford 2004) and so on.

The sensing process can be divided in two main areas; (a) monitoring of the user and his/her activities, and, (b) monitoring of the environment itself.

User's behaviour monitoring, tracking and identification of people in an environment is an important issue in AmI systems. Motion sensors have been used in many applications to track users. However, while they can sense movement they cannot identify the user. For that, new alternatives have been produced allowing to track and identify the users (and objects). One of the main examples of these new technologies are RFID tags that can be coupled with RFID readers to monitor the movement of the users and objects (Schneider et al. 2009). It is worth mentioning that all these methods have limitations and cannot guarantee accurate identification in all cases however they provide interesting data that can be used for the tracking and identification.

Monitoring the state of environment is as important as monitoring the user. Knowing, for example, the temperature, the humidity and the light level of a room allows a system to contextualize the actions of the user. Sensors capable to measure temperature, humidity and luminosity are being used in various different projects.

Sensing plays then two important roles in AmI systems. On the one hand it provides data to learn how the user behaves, and on the other hand it provides information about the current state of the environment in order to make correct decisions at each moment.

2.2 Reasoning

The technologies included in the reasoning step are the key in the process of providing the environment with intelligence. Figure 2 illustrates the relationships among these technologies.

Acting intelligently demands to recognise the action/activity the user is carrying out at each moment as well as the general action/plan the person is performing. The approaches to consider differ according to the type of sensors and the set of activities to recognize. For example in order to recognize repetitive body motions, e.g., walking, data collected from accelerometers is very useful (Maurer et al. 2006). If the action is not as easily distinguishable by body position, adding sensors in daily object such as doors, kettle or medicine dispenser (Tapia et al. 2004), can be a possible solution. In addition, some researchers (Brdiczka et al. 2007) use video cameras to identify activities. Regarding to machine learning techniques that have been used for activity recognition, Nave Bayes classifiers have shown promising results (Doctor et al. 2005; Tapia et al. 2004; Brdiczka et al. 2007; Muehlenbrock et al. 2004).

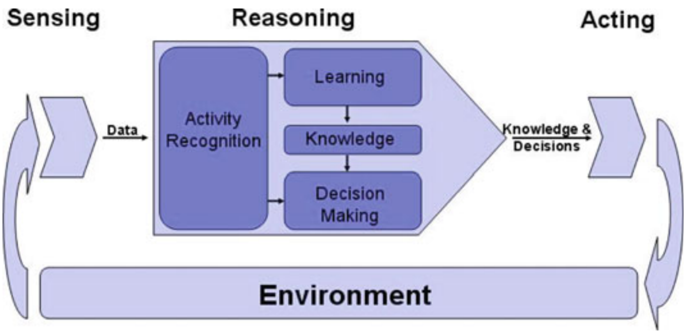

Fig. 2 Reasoning technologies

Once the behaviours of the users have been learnt and the activities the user is doing have been recognized, it is the moment of making decisions to act proactively. This area has attracted little attention among AmI researchers so far. As notable exception we can highlight the analysis of Simpson et al. (2006) about how AI planning systems can be used in AmI decision making systems and the rulebased system suggested by Augusto and Nugent (2004). A more integrated approach mixing the learning and decision making processes (Mozer 2004) (See Sect. 4.1).

As AmI systems are related to real world environments, very little can be done without taking into account the spatial and temporal aspects. Gottfried et al. (2006) analyse how traditional frameworks for spatial and temporal reasoning can be used in AmI systems. Aztiria et al. (2008a) have also related the spatial and temporal aspects to AmI environments. They define a language, which is strongly based on spatial and temporal information, that allows us to represent user's behaviours and preferences in a clear and non-ambiguous format. The importance of spatial and temporal aspects in the learning process will be explored in Sect. 3.

2.3 Acting

Intelligent environments can act over the environment in a variety of ways. Following the decision making process the system may decide that some conditions needs adjustment (for example, closing the curtains or starting the heating) or interaction with humans is needed (for example, to alert a carer if a person in the house has fallen). The most natural way of acting is by means of actuators integrated in standard devices. The use of friendly multimodal user interfaces that aim at making the interaction as close to human-to-human communication as possible (e.g., through voice and image processing technologies) can also be considered (Turunen et al. 2007; Partala et al. 2006; Coen 1998). Robots are also able to assist people in different tasks (Pineau et al. 2003) and they can even show emotions closer to human-beings (Brezeal and Scassellati 1999).

Acting necessarily involves users so that these aspects are considered in the learning process.

3 AmI's special features

Once the necessary technologies to create a real AmI environment have been defined, we will focus on the learning process. In that sense, some AmI features must be specially taken into account in the learning process. These features can be divided into four groups.

– The importance of the user (A)
– Data captured by sensors (B)
– Acquired patterns and their representation (C)
– Timing the learning process (D)

In the following subsections these characteristics have been divided in sub-characteristics (A.I, B.I, B.II,...) in order to define more precise relationships of different ML techniques with AmI systems in Sect. 4.

3.1 Users in AmI environments

Users are the focus of any development in AmI, and the fact that the environment is technologically rich must not translate into any extra effort on behalf of the users to obtain benefits of an AmI system (Dooley et al. 2006; Muller 2004). This implies that the data acquisition and feedback obtaining processes have to be carried out as unobtrusively as possible (A.I) (Rutishauser et al. 2005) (e.g., by means of sensors installed in standard devices). As explained in Sect. 2.3, other ways of interaction are feasible although a delicate and subtle balance has to be achieved by an AmI system such that it helps without being overwhelming (Liao et al. 2004; Pollack 2005). Importunate systems will decrease the chances of massive technology adoption.

3.2 Data captured by sensors

The importance of sensing (See Sect. 2.1) increases when considering the learning process. The data collected from the sensors will greatly influence the learning process, and all patterns will depend upon the data captured. This dependency is hindered by technical problems associated with the gathering and interpretation of the data sensors produce.

First, data will be collected in a continuous way from different information sources. AmI systems may employ a centralized or distributed model (Akyldiz et al. 2002). Integrating data from different sources usually presents many challenges because different sources will use different styles of record keeping and different storage modes (e.g., digital vs. analog). Thus, data fusion techniques (such as Kalman filters) (Brammer and Siffling 1989) as well as probabilistic approaches to combine information will be necessary (Berger 1985; Manyika and Durrant-Whyte 1994). Handling temporal information will be one of the most challenging aspects. Moreover, as in other areas, noisy data with missing or inaccurate values will be common. Determining how to appropriately deal with these aspects is a challenge for consideration (B.I).

It is also necessary to consider the nature of the collected data (B.II). The first aspect to be considered is the nature of raw data. Sometimes, raw data will not be meaningful enough, and a combination of different sources of raw data will be necessary in order to infer and recognize meaningful activities (as explained in Sect. 2.2). For example, in order to infer that 'Michael has gone into the bathroom', we will have to combine the raw data of 'There is a motion in the corridor', 'RFID in the bathroom door detects that Michael is passing through the door' and 'There is motion inside the bathroom'. In that sense, 'activity recognition' is becoming an independent and important field in AmI (Cook et al. 2009).

Another aspect to consider is that different types of sensors provide information of different natures that can be used for different purposes in the learning process (B.III). Some sensors provide direct information about the actions of the user (e.g., a sensor installed in the bathroom's light switch provides direct information about when someone switched on the light). Other sensors provide information about the environment itself (e.g., a temperature sensor installed in the bathroom). Finally, we can also consider sensors that provide information about the health and emotional status of the user (e.g., sensors that capture parameters like heart rate).

Finally, due to the complexity of AmI environments, externally gathered information can be considered in order to enrich the collected data (B.IV). Externally gathered information will typically be domain-related knowledge, such as the like medical background of patients, preferences of the user specified in advance by the user or calendar information (e.g., when the user goes on holiday).

3.3 Acquired patterns and their representation

As explained above, the learned patterns can be used for many different purposes. Thus, they sometimes have to be integrated into a bigger system in order to make sensible high-level decisions. Therefore, the learning system must take into account requirements of a wider range of systems. If learned patterns are used in order to understand behaviors of the user or combined with other types of knowledge to obtain more meaningful information, such processes will normally demand comprehensible patterns. Moreover, it may be necessary for the system to explain to the user the decisions made. Therefore it is essential that the output and process be easily transformed into a human-understandable language (Leake et al. 2006) (C.I).

If a comprehensible representation is demanded, it can be obtained in different ways. So far, the most common representation relates the action of the user to the status of any type of sensor (Hagras et al. 2004; Gal et al. 2001). In that sense, some authors (Duman et al. 2008; Jakkula et al. 2007) have seen the need to relate user actions in order to facilitate the understanding of the user's behavior (C.II). Representing the behaviors relating actions (and even representing the behavior as a sequence of actions) allows us to use relative time references instead of absolute times. Such time relations could be qualitative or quantitative time relations (C.III). Finally, in addition to the need to specify sequenced actions and their time relations, the need to use conditions in order to contextualize the sequence (e.g., usually from 7:50 onwards) and know how behavior varies from time to time (e.g., on weekdays compared to weekends) is clear (C.IV). Considering all of these aspects, Michael's morning habits can be represented as shown in Fig. 3.

3.4 Scheduling the learning processes

Discovering common patterns of behavior in recorded data is part of a complex process involving several stages. Acting as intelligently as possible while the system is gathering

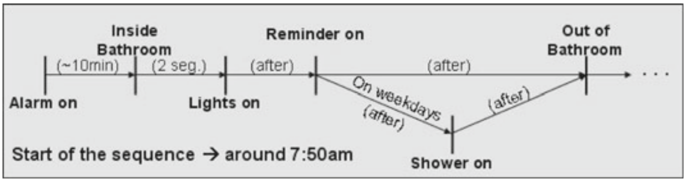

Fig. 3 Sequence of Michael's morning habits

data for the learning process would be convenient (D.I). Although these actions will not be as intelligent and efficient as those performed once the patterns of the user have been learned, minimum services can be provided without such patterns. As we will see in Sect. 4, some ML techniques can effectively carry out this initial process. Even the use of stereotypes has been suggested by some research to be useful to expedite the initial learning process (Kofod-Petersen 2006).

On the other hand, once well-defined patterns have been extracted from the data collected through the sensors, those patterns have to be continuously revised and updated (D.II) because

- The user can change his/her preferences (e.g., when summer arrives, Michael prefers having a shower that is not as hot as in winter).
- New patterns could appear (e.g., Michael is becoming interested in financial news, so he wants to watch financial news as well as news about his favorite soccer team in the morning).
- Previously learnt patterns were incorrect (e.g., every day Michael has to switch on the light in the bathroom because the system switches it on 10 s after he goes into the bathroom).

This adaptation process could mean the modification of parameters in a previously learned pattern, addition of a new pattern or even deletion of some previously learnt pattern. The adaptation period lasts for as long as the intelligent environment is in operation. Hence, collecting user feedback either by means of standard devices or by means of other systems is necessary.

We conclude that at least three main stages can be identified: (i) a stage in which the system act as intelligently as possible without patterns while the system is gathering data, (ii) the main learning process stage, where user's patterns are extracted from collected data and (iii) a stage of periodic revision and adaptation of patterns to reflect meaningful changes in a dynamic environment.

4 Applications and machine learning techniques

It should not come as a surprise that no research on learning for AmI reported in the literature covers all of the aspects considered in the previous section. After analyzing applications developed by different groups, we realized that current applications are very specific applications with focused goals. Therefore, they use the most suitable ML technique and only consider the needs of a specific application. In addition to analyzing what knowledge is being learned in each application, we took as a guide the ML technique employed to analyze the strong and weak aspects of each technique. Additionally, we considered the special features of AmI environments mentioned in Sect. 3 (A.I, B.I, B.II, ...).

Due to space restrictions, we only considered those applications that attempt to discover patterns of user behavior by means of ML techniques. We have thus avoided areas of AmI like activity recognition or anomaly detection, where ML techniques have also been used.

4.1 Artificial neural networks (ANNs)

Mozer et al. (1995) and Chan et al. (1995) were amongst the first reports on applications for AmI environments in which user patterns were considered.

In the case of Mozer et al., the aim of their system (installed in the Adaptive House) was to design an adaptive control system that considers the lifestyle and energy consumption of the inhabitants. For that, they used a feed-forward neural network to develop an occupancy predictor and zone anticipator, which were used to predict where the user would be in the coming seconds and control lighting based on those predictions. Chan et al. developed a similar application in order to asses if a given situation was normal or abnormal. After validating this application in an institution for elderly and disabled people, they claimed that 9 out 10 predictions by the system were correct.

Other authors have kept using ANNs in order to provide personalized services. Thus, Campo et al. (2006) developed a system that calculated the probability of occupation for each area of the house and systematically compared the probability with the current situation. The intelligent thermostats of Boisvert and Rubio (1999), which reduced energy consumption by 9–16%, also employed ANNs. See Begg and Hassan (2006) for a survey focused on ANNs for Smart Homes.

We next considered the use of ANNs in the scenario suggested in the Introduction. Due to their capacity to manage complex data and create complex models (B.I), we assert that the system would provide correct responses in situations such as turning on the lights when Michael goes into the bathroom or getting ready the shower on weekdays. However, the use of ANNs has a limitation related to their black box nature; therefore their internal structure is not human-readable (C.I). Thus, the system would be able to turn on the light but would not be able to explain, in a comprehensible way, how it inferred such an output. If we were looking for user patterns in order to understand common behaviors of the user, the system would not be able to represent the learned patterns in a comprehensible way. In AmI environments where user plays a central role, the ability to represent patterns and explain the actions carried out by the environment is essential.

4.2 Classification techniques

The group that works in the environment named 'SmartOffice' (Gal et al. 2001) pointed out that 'a user is only willing to accept an intelligent environment offering services implicitly if he understands and foresees its decisions'. Brdiczka et al. (2005). Thus, this group ruled out ANNs and, decided to use rules to represent patterns of the user. In particular, they used decision trees, which can be translated into rules (C.I). Their algorithm is based on splitting situations in which examples indicate different reactions. For Michael's example, analyzing what Michael does when he goes into bathroom and turns on the light, allows the system to realize that sometimes he takes a shower and sometimes he does not. In these situations, the algorithm tries to discover the conditions for each case. In this case, it would discover that he takes a shower on weekdays but not weekends (C.IV).

Other authors (Stankovski and Tmkoczy 2006) have also pointed out that the model generated by decision trees is easy to understand and suitable for human inspection. In this case, they suggested using a decision tree in order to detect abnormal situations. For that, they trained a decision tree with the collected data that described the normal state of the environment. Thus, any situation external to the tree would be considered as abnormal. In Michael's case, it could detect hazardous situations (e.g., when Michael leaves the cooker on when leaving the house). A system constructed under such hypothesis can generate many "false positives", however, because new situations inherent to AmI environments are not always abnormal.

Rules are another knowledge representation mechanism often used in AmI systems. In addition to their human-readable representation, they have the advantage that they are easy to add, modify or delete (D.II). Thus, if Michael changes his preferences and decides to take a shower only on Mondays, Wednesdays and Saturdays instead all weekdays, it would be very easy to update this knowledge. On the other hand, rules do not relate actions across several rules but instead relate a specific situation to an action, giving no sense of sequence to the actions (C.II).

4.3 Fuzzy logic rules

Researchers at Essex's iDorm lab have given prominence to the problem of learning and are one of the most active groups in this area. Their initial efforts (Hagras et al. 2004; Doctor et al. 2005) were focused on developing an application that generated a set of fuzzy rules representing a user's patterns. Recording the changes caused by the user in the environment, they generated membership functions as well as fuzzy rules that mapped data into fuzzy rules (C.I). Additionally, they defined a strategy to adapt the rules based on negative feedback given by the user (D.II). They validated this approach at iDorm over the course of 5 days using seven input sensors and ten output actuators, generating 308 rules (from 408 instances recorded).

The nature of rules generated in this way will be in some general sense similar to that of rules obtained through the classification techniques described in the previous subsection. They are considered more robust when dealing with data of a continuous nature (e.g., temperature, humidity and time) (B.I). Due to the number of sensors and different situations that can be generated when combining sensors, it seems clear that relating actions only to the global situation (without relating actions among them) will result in an excessive number of generated rules with very little meaning (C.II).

Thus, an improvement (Duman et al. 2008) involved identifying relevant and important associations between given actions, so that irrelevant aspects of the rules (and, by extension, some rules as well) could be removed. In Michael's case, it is clear that the action of turning on the lights in the bathroom is typically associated with the action of going into the bathroom; after statistical analysis, therefore, only frequent associations above an established threshold would remain. Although relations defined in such a way are not complete sequences of actions but instead one-to-one relations, they relate actions to actions instead of relating actions to the status of all sensors. In a experiment carried out in the same environment, this allowed for a 91% reduction in the number of rules.

Vainio et al. (2008) also used fuzzy rules to represent habits of the user. In contrast to the approach followed in the iDorm project, these authors manually constructed the membership functions and used reinforcement learning to replace old rules in order to prevent single overriding event from having a large impact (D.I).

4.4 Sequence discovery

The group that has been working on the MavHome and Casas projects is one of the most active groups in this area. The first applications developed by this group were focused on building universal models, represented by means of Markov models, in order to predict either future locations or activities (Cook and Das 2007). In this sense, they made notable improvements by developing applications to discover daily and weekly patterns (Heierman and Cook 2002). Additionally, they constructed application with the ability to infer abstract tasks automatically, identifying corresponding activities that were likely to be part of the same task (Rao and Cook 2004).

One of the major contributions of researchers from this group has been the use of time intervals associated with activities (Jakkula et al. 2007) rather than instantaneous time points (C.III). They used Allen's temporal logic (Allen 1984) in order to represent time intervals, producing fairly intuitive sequences of actions. Considering again Michael's example, the system would be able to detect that he first gets up, then goes into the bathroom and then turns on the lights.

A few aspects requiring improvement can be noted. First, this system does not detect all activities as part of the same sequence but instead detects relations separately (relating actions in a one-to-one manner). Second, this system only considers Allen's temporal logic relations (which define relations qualitatively), ruling out quantitative relations (C.III). Thus, the term 'after' means that Michael turns on the lights when he goes into the bathroom; however, the likely delay between one action and the next cannot be measured. Defining relations by means of quantitative values allows the system to automate actions, which is impossible with purely qualitative values (e.g., the system knows that turning the lights on comes after a given event, but it does not know if is the time delay is 2 s, 5 min or 2 h after the first event). Finally, it is worth mentioning that this method does no discover conditions; such a concept is in fact very useful if two activities are related in different ways. In Michael's case, the relation 'turn on the light, after, have a shower' is real only on weekdays so that condition should be considered in order to provide more accurate patterns. Experiments carried out using a prediction technique improved when using temporal rules. In particular, prediction performance improved in 1.86% of tests carried out with real data and 7.81% of tests with synthetic data.

4.5 Instance-based learning

Research conducted under the MyCampus project filtered messages based on the preferences of the user (Sadeh et al. 2005). For that, this project used Case-Based Reasoning (CBR), which can be defined as an Instance-Based Learning (IBL) technique. First, it tried to filter messages based on a limited set of a priori preference options. However, only 50% of participants indicated they were only satisfied with the messages they had received using the system. For that reason, they added a simple CBR module that matched cases using a Nearest-Neighborhood algorithm. This significantly improved the quality of filtering (i.e., user satisfaction increased from 50 to 80%). UT-AGENT (Kushwaha et al. 2004) also used CBR in order to learn the preferences of the user. By recording the set of tasks that the user usually performed, the UT-AGENT tried to provide information or help to the user.

Let us consider the use of IBL techniques in Michael's example. Given a situation that is similar to previous ones, the system would act properly because IBL techniques provide similar solutions to similar problems/situations without any initial model (D.I). Thus, when Michael goes into the bathroom, the system would compare that situation to previous ones and correctly turn on the lights. The same would occur for other situations similar to previous occurrences (e.g., having a shower on weekdays or maintaining the temperature of the house around 24–26°C).

However, use of IBL techniques has some limitations. As this process infers a solution for each specific situation, it does not create a model that represents patterns (C.I). Therefore, it would not be possible to extract a general pattern indicating that Michael turns on the lights after going into the bathroom. Further, as each situation can be represented by means of a large number of parameters, the matching process could be very difficult because there are no clues regarding the importance of each parameter in each situation (B.III). Thus, if we consider the parameter 'day of the week' when Michael turns on the lights, it seems clear when he takes a shower and when he does not. However, other parameters (e.g., light level, temperature) that would shape the pattern in a different way could also be considered.

Finally, it is worth mentioning that new approaches related to IBL techniques, such as temporal CBR (Galushka et al. 2006), are emerging.

4.6 Reinforcement learning

Research conducted under the Adaptive House and SmartOffice projects has used a module based on reinforcement learning in order to add the capacity of adaptation. Mozer (2004) used Q learning for lighting regulation (on/off status and intensity). Taking as a starting point the fact that the inhabitant had no preference for the device setting, the system tried to minimize energy consumption as long as the inhabitant did not express discomfort. Once the system received feedback from the user, it tried to balance user preferences with energy consumption. In research from the Smart Office project reported in Zaidenberg et al. (2008), researchers tried to adapt a pre-defined set of actions progressively to particular users by giving the system rewards associated with good decisions.

In Michael's example, if we consider that the system already has a model (either defined manually or learned by means of previously mentioned techniques), reinforcement learning techniques can be used in order to adapt such patterns (D.II). Let us hypothesize that learned patterns define that the shower temperature must be around 24–26°C. Every time Michael takes a shower, however, he changes the settings to indicate a preference of 28–29°C. This would be considered negative feedback. After collecting feedback, reinforcement learning would change the pattern and adapt it to Michael's new preferences.

Still, the use of this technique demands a set of initial patterns that ideally should be learned automatically instead of from pre-defined models (which could annoy users and even make the process of learning habits without any bias difficult). Although other techniques have the same limitation, the inherent difficulty in reinforcement learning is interpreting user's feedbacks; this is particularly important for reinforcement learning because this system is based mainly on the interpretation of this feedback (A.I).

4.7 Summary of learning techniques

As we have seen in the previous sections, each technique has strengths and weaknesses. In order to summarize these aspects, Tables 1 and 2 highlight the strengths and weaknesses of each technique relative to AmI environments characteristics and the different learning stages defined in Sect. 3.4.

Analysis of the different techniques reveals that different groups have chosen techniques based on the specific needs of their environments or applications. We assert that still there is no global or holistic approach that allows the system to learn easy-to-understand patterns, represent human behavior by means of sequences of actions, relate actions to actions, discover conditions, act intelligently while recording data or adapt patterns.

Given the strengths and weaknesses of each technique, combining different techniques seems to offer a promising approach. IBL techniques (see Sect. 4.5) can be considered to allow the system to act intelligently while the system is recording data. Once data is recorded, techniques to discover sequences (see Sect. 4.4) can be used in order to identify the user's common behaviors and define time relations using either Allen's temporal logic and/or quantitative relations. In order to discover accurate conditions, classification techniques can be useful (see Sect. 4.2). Finally, reinforcement learning techniques (see Sect. 4.6) can be used to adapt the learned patterns. Some authors (Aztiria et al. 2008b) have started to use this strategy of combining different techniques in order to identify complex patterns in AmI environments.

Table 1 Strengths and weaknesses of learning techniques (part 1)

	Strengths/weaknesses	Artificial neural network	classification techniques	Fuzzy logic rules
Short period learning without patterns	Strengths	–	–	–
	Weaknesses	Need of data for training (D.I). No human-readable output (C.I)	Need of data for training (D.I)	Need of data for training (D.I)
Pattern extraction process from collected data	Strengths	Capacity to generalize and model complex situations (B.I)	Human-readable output (C.I). Discovering of conditions (C.IV)	Human-readable output (C.I). Robust to uncertainties (B.I).
	Weaknesses	No human-readable output (C.I)	Event-Situation relations only (C.II)	Event-Situation relations only (C.II)
Adaptation process to reflect changes in dynamic environments	Strengths	Possibility of introducing neurons dynamically (D.II)	Easy to add, delete or change rules (D.II)	Easy to add, delete or change rules (D.II)
	Weaknesses	Need of restructuring of the network (C.I)	Need of restructuring to avoid more conflicts (C.I)	Need of restructuring to avoid conflicts (C.I)

5 Conclusions and ongoing challenges

Until now, research related to AmI has devoted significant attention to supporting technologies (e.g., physical components and middleware). This survey emphasizes the importance of advances in Machine Learning for fully realizing the AmI paradigm. We provided evidence that an increasing effort is being made to adapt learning techniques to the development of AmI environments. We highlighted the need for further work because the focus so far has been on applying well-known Machine Learning techniques rather than developing new techniques specifically for AmI environments.

Muller (2004) pointed out that 'In many research projects, great results were achieved …but the overall dilemma remains: there does not seem to be a system that learns quickly, is highly accurate, is nearly domain independent, does this from few examples with literally no bias, and delivers a user model that is understandable and contains breaking news about the characteristics of the user. Each single problem favours a certain learning approach'. We suggest that we are still in a situation in which a holistic approach has not been achieved. Being aware of the needs of each particular environment will allow us to identify different complementary techniques that can fulfill our needs when suitably assembled. Even so, we

Table 2 Strengths and weaknesses of learning techniques (part 2)

	Strengths/weaknesses	Sequence discovery	Instance based learning	Reinforcement learning
Short period learning without patterns	Strengths	–	No need for training. Reasoning based upon previous cases (D.I)	–
	Weaknesses	Need of data for training (D.I)	Sensitive to large attributes set and importance of each one (B.III)	Need for a model
Pattern extraction process from collected data	Strengths	Rules that represent relations among events (C.II). Discovering qualitative time relations (C.III)	Possible source of information for other techniques	–
	Weaknesses	Algorithms do not focus on discovering quantitative time intervals (C.III)	No creation of general patterns (C.I)	Need for a model
Adaptation process to reflect changes in dynamic environ-ments	Strengths	Easy to add, delete or change sequences (D.II)	Possibility of providing a solution to instances outside of patterns (D.II)	Adapts patterns after interacting with the user (D.II)
	Weaknesses	Possible conflicts when adding new sequences (C.I)	Computational and temporal cost	Difficulties understand-ing feedback (A.I)

think that the design and development of a learning technique suitable for all environments would be ideal.

As this is a new development within computer science, learning in AmI environments is full of challenges to be overcome. Ongoing challenges in establishing a suitable learning paradigm for AmI are

- Being able to handle environments where multiple inhabitants co-exist.
- Being able to accurately infer high-level tasks from the low-level information.
- Being able to manage the complexity and richness of everyday human life.

We also want to stress the importance of security and privacy in these environments. Behavioral patterns (i.e., when the user usually leaves home and when he/she will return) are usually considered personal and sensitive. These should therefore be stored in a secure way to prevent public access as well as preclude potentially disastrous consequences (e.g., a criminal accessing the information).

Although the perfect learning system for AmI environments has not been achieved yet, we conclude that each contribution to this area brings us one step closer to the realization of Ambient Intelligence. This progress allows the AmI systems to understand what the user wants and needs as well as also where, when and to whom to deliver a service.

Acknowledgments The authors are grateful to Julie-Ann Walkden (from the South Eastern Trust) and the anonymous reviewers who provided helpful suggestions to improve the content of this article.

References

Aarts E (2004) Ambient intelligence: a multimedia perspective. IEEE Multimed 11(1):12–19

Akyldiz IF, Su W, Sankarasubramanian Y, Cayirci E (2002) A survey on sensor networks. IEEE Commun Mag 40:102–114

Allen J (1984) Towards a general theory of action and time. In: Artificial Intelligence, vol. 23. pp 123–154

Augusto J (2007) Ambient intelligence: the confluence of pervasive computing and artificial intelligence. In: Schuster A (ed) Intelligent computing everywhere. Springer, Berlin, pp 213–234

Augusto J, Cook D (2007) Ambient Intelligence: applications in society and opportunities for AI. In: Tutorial lecture notes delivered at 20th international joint conference on artificial intelligence (IJCAI-07). IJCAI, Hyderabad, India

Augusto JC, Nugent CD (2004) The use of temporal reasoning and management of complex events in smart homes. In: Proccedings of European conference on AI (ECAI 2004). IO Press, pp 778–782

Augusto JC, Nugent CD (2006) Smart homes can be smarter. In: Augusto JC, Nugent CD (eds) Designing smart homes. The role of artificial intelligence. Springer, Berlin, pp 1–15

Aztiria A, Augusto JC, Izaguirre A (2008a) Spatial and temporal aspects for pattern representation and discovery in intelligent environments. In: Workshop on spatial and temporal reasoning at 18th European conference on artificial intelligence (ECAI 2008) (to be published)

Aztiria A, Augusto JC, Izaguirre A, Cook DJ (2008b) Learning accurate temporal relations from user actions in intelligent environments. In: Proceedings of the 3rd symposium of ubiquitous computing and ambient intelligence, vol. 51/2009. pp 274–283

Begg R, Hassan R (2006) Artificial neural networks in smart homes. In: Augusto JC, Nugent CD (eds) Designing smart homes. The role of artificial intelligence. Springer, Berlin, pp 146–164

Berger JO (1985) Statistical decisions. Springer, Berlin

Boisvert A, Rubio RB (1999) Architecture for intelligent thermostats that learn from occupants' behavior. In: ASHRAE transactions. pp 124–130

Brammer K, Siffling G (1989) Kalman bucy filters. Artech House

Brdiczka O, Reignier P, Crowley JL (2005) Supervised learning of an abstract context model for an intelligent environment. In: Proceedings of the 2005 joint conference on smart objects and ambient intelligence: innovative context-aware services: usages and technologies, vol. 121. ACM, pp 259–264

Brdiczka O, Reignier P, Crowley J (2007) Detecting individual activities from video in a smart home. In: Proceedings of the international conference on knowledge-based and intelligent information and engineering systems. pp 363–370

Brezeal C, Scassellati B (1999) A context-dependent attention system for a social robot. In: Proceedings of the international joint conference on artificial intelligence. pp 1146–1151

Campo E, Bonhomme S, Chan M, Esteve D (2006) Learning life habits and practices: an issue to the smart home. In: international conference on smart Homes and health telematic. pp 355–358

Chan M, Hariton C, Ringeard P, Campo E (1995) Smart house automation system for the elderly and the disabled. In: Proceedings of the 1995 IEEE international conference on systems, man and cybernetics. pp 1586–1589

Coen MH (1998) Design principles for intelligent environments. In: Proceedings of the 1998 15th national conference on artificial intelligence, AAAI. AAAI Press, pp 547–554

Cook D, Augusto J, Jakkula V (2009) Ambient intelligence: applications in society and opportunities for ai. Pervasive Mob Comput 5(4):277–298

Cook DJ, Das SK (2005) Smart environments: technology, protocols and applications. Wiley, New York

Cook DJ, Das SK (2007) How smart are our environments? an updated look at the state of the art. In: Pervasive and mobile computing, Elsevier science, vol. 3. pp 53–73

Cook DJ, Huber M, Gopalratnam K, Youngblood M (2003) Learning to control a smart home environment. In: Innovative applications of artificial intelligence

Das SK, Cook DJ (2006) Designing and modeling smart environments. In: Proceedings of international sysmposium on wireless, mobile and multimedia networks. pp 490–494

Delapierre G, Grange H, Chambaz B, Destannes L (1983) Polymer-based capacitive humidity sensor. Sens Actuators 4(1):97–104

Doctor F, Hagras H, Callaghan V (2005) A fuzzy embedded agent-based approach for realizing ambient intelligence in intelligent inhabited environments. In: IEEE transactions on systems, man and cybernetics: part A, vol. 35. pp 55–65

Dooley J, Callaghan V, Hagras H, Bull P, Rohlfing D (2006) Ambient intelligence—knowledge representation, processing and distribution in intelligent inhabited environments. In: 2nd IET international conference on intelligent environments, IE 06, pp 51–59

Ducatel K, Bogdanowicz M, Scapolo F, Leijten J, Burgelman J (2001) Scenarios for ambient intelligence in 2010. Technical report URL http://cordis.europa.eu/ist/istagreports.htm

Duman H, Hagras H, Callaghan V (2008) Intelligent association exploration and exploitation of fuzzy agents in ambient intelligent environments. J Uncertain Syst 2(2):133–143

Ermes M, Parkka J, Mantyjarvi J, Korhonen I (2008) Detection of daily activities and sports with wearable sensors in controlled and uncontrolled contitions. IEEE Trans Inf Technol Biomed 12:20–26

Friedemann M, Mahmoud N (2002) Pervasive computing, first international conference. Springer

Friedwald M, Costa OMD (2003) Science and technology roadmapping: ambient intelligence in everyday life (ami@life)

Friedwald M, Costa OMD, Punie Y, Alahuhta P, Heinonen S (2005) Perspectives of ambient intelligence in the home environment. In: Telematics and informatics, vol. 22. Pergamon Press, pp 221–238

Gal CL, Martin J, Lux A, Crowley JL (2001) Smartoffice: design of an intelligent environment. IEEE Intell Syst 16(4):60–66

Galushka M, Patterson D, Rooney N (2006) Temporal data mining for smart homes. In: Augusto JC, Nugent CD (eds) Designing smart homes. The role of artificial intelligence. Springer, Berlin, pp 85–108

Gottfried B, Guesgen HW, Hubner S (2006) Spatiotemporal reasoning for smart homes. In: Designing smart homes. The role of artificial intelligence, m1: copyright 2006, the institution of engineering and technology. Springer, pp 16–34

Hagras H, Callaghan V, Colley M, Clarke G, Pounds-Cornish A, Duman H (2004) Creating an ambient-intelligence environment using embedded agents. IEEE Intell Syst 19(6):12–20

Heierman EO, Cook DJ (2002) Improving home automation by discovering regularly occurring device usage patterns. In: Third IEEE international conference on data mining. pp 537–540

Jakkula VR, Crandall AS, Cook DJ (2007) Knowledge discovery in entity based smart environment resident data using temporal relation based data mining. In: 7th IEEE international conference on datamining. pp 625–630

Jiang L, Liu DY, Yang B (2004) Smart home research. In: Proceedings of 2004 international conference on machine learning and cybernetics, vol. 2. pp 659–663

Kautz H, Fox D, Etzioni O, Borriello G, Arnstein L (2002) An overview of the assisted cognition project. In: Proceedings of the AAAI workshop on automation as Caregiver. AAAI Press, pp 60–65

Kofod-Petersen A (2006) Challenges in case-based reasoning for context awareness in ambient intelligence systems. In: 1st workshop on case-based reasoning and context awareness

Kushwaha N, Kim M, Kim DY, Cho W (2004) An intelligent agent for ubiquitous computing environments: smart home ut-agent. In: Proceedings of the 2nd IEEE workshop on software technologies for future embedded and ubiquitous systems. pp 157–159

Leake D, Maguitman A, Reichherzer T (2006) Cases, context, and comfort: opportunities for case-based reasoning in smart homes. In: Augusto JC, Nugent CD (eds) Designing smart homes. The role of artificial intelligence. Springer, Berlin, pp 109–131

Liao L, Patterson DJ, Fox D, Kautz H (2004) Behavior recognition in assisted cognition. In: Proceedings of the IAAA-04 workshop on supervisory control of learning and adaptive systems. pp 41–42

Manyika J, Durrant-Whyte H (1994) Data fusion and sensor management: a decentralized information-theoretic approach. Ellis Horwood, Chichester

Maurer U, Smailagic A, Siewiorek D, Deisher M (2006) Activity recognition and monitoring using multiple sensors on different body positions. In: Proceedings of the international workshop on wearable and implantable body sensor networks. pp 99–102

Mitchell TM (1997) Machine learning. McGraw-Hill/MIT Press

Mozer MC (2004) Lessons from an adaptive home. Smart environments: technology, protocols and applications. Wiley, pp 273–298

Mozer MC, Dodier RH, Anderson M, Vidmar L, Cruickshank RF, Miller D (1995) The neural network house: an overview. Erlbaum, pp 371–380. Current trends in connectionism

Muehlenbrock M, Brdiczka O, Snowdon D, Meunier J (2004) Learning to detect user activity and availability from a variety of sensor data. In: Proceedings of the IEEE international conference on pervasive computing and communications

Muller ME (2004) Can user models be learned at all? Inherent problems in machine learning for user modelling. In: Knowledge engineering review, vol. 19. Cambridge University Press, pp 61–88

Pantic M, Pentland A, Nijholt A, Huang T (2006) Human computing and machine understanding of human behavior: a survey. In: Proceedings of the 8th international conference on multimodal interfaces. ACM, pp 239–248

Partala T, Surakka V, Vanhala T (2006) Real-time estimation of emotional experiences from facial expressions. Interact Comput 18(2):208–226

Pineau J, Montemerlo M, Pollack M, Roy N, Thrun S (2003) Towards robotic assistants in nursing homes: challenges and results. Rob Auton Syst 42(3–4):271–281

Pollack ME (2005) Intelligent technology for an aging population: the use of ai to assist elders with cognitive impairment. AI Magazine 26(2):9–24

Ramos C, Augusto J, Shapiro D (2008) Ambient intelligence—the next step for artificial intelligence (guest editors' introduction to the special issue on ambient intelligence). IEEE Intell Syst 23(2):15–18

Rao SP, Cook DJ (2004) Predicting inhabitant action using action and task models with application to smart homes. Int J Artif Intell Tools (Architectures, Languages, Algorithms) 13(1):81–99

Rivera-Illingworth F, Callaghan V, Hagras H (2005) A neural network agent based approach to activity detection in ami environments. In: IEEE international workshop on intelligent environments. pp 92–99

Rutishauser U, Joller J, Douglas R (2005) Control and learning of ambience by an intelligent building. In: IEEE on systems, man and cybernetics: a special issue on ambient intelligence, IEEE systems, man, and cybernetics society. pp 121–132

Sadeh NM, Gandom FL, Kwon OB (2005) Ambient intelligence: the mycampus experience. Technical report CMU-ISRI-05-123, ISRI

Schneider M, Krner A, Alvarado JCE, Higuera AG, Augusto JC, Cook DJ, Ikonen V, Cech P, Mikuleck P, Kameas A, Callaghan V (2009) 1st international workshop on rfid technology: concepts, practices and solutions. included in workshops proceedings of the 5th international conference on intelligent environments

Shadbolt N (2003) Ambient intelligence. IEEE Intell Syst 18(4):2–3

Simpson R, Schreckenghost D, LoPresti EF, Kirsch N (2006) Plans and planning in smart homes. In: Augusto J, Nugent C (eds) Designing smart homes: the role of artificial intelligence. Springer, Berlin

Stanford V (2004) Biosignals offer potential for direct interfaces and health monitoring. IEEE Pervasive Comput 3:99–103

Stankovski V, Tmkoczy J (2006) Application of decision trees to smart homes. In: Augusto JC, Nugent CD (ed) Designing Smart Homes. The role of artificial intelligence, m1: copyright 2006, the institution of engineering and technology. Springer, pp 132–45

Tapia EM, Intille SS, Larson K (2004) Activity recognition in the home using simple and ubiquitous sensors. In: Proceedings of pervasive. pp 158–175

Turunen M, Hakulinen J, Kainulainen A, Melto A, Hurtig T (2007) Design of a rich multimodal interface for mobile spoken route guidance. In: Proceedings of interspeech 2007—Eurospeech

Vainio AM, Valtonen M, Vanhala J (2008) Proactive fuzzy control and adaptation methods for smart homes. IEEE Intell Syst 23(2):42–49

Weiser M (1991) The computer for the 21st century. Sci Am 265(3):94–104

Witten IH, Frank E (2005) Data mining: practical machine learning tools and techniques, 2nd edn. Elsevier, Amsterdam

Wolffenbuttel RF, Mahmoud KM, Regtien PL (1990) Compliant capacitive wrist sensor for use in industrial robots. IEEE Trans Instrum Meas 39:991–997

Zaidenberg S, Reignier P, Crowley JL (2008) Reinforcement learning of context models for a ubiquitous personal assistant. In: Proceedings of the 3rd symposium of ubiquitous computing and ambient intelligence, vol. 51/2009. pp 254–264