Ambient Intelligence: A Survey

FARIBA SADRI, Imperial College

36

In this article we survey ambient intelligence (AmI), including its applications, some of the technologies it uses, and its social and ethical implications. The applications include AmI at home, care of the elderly, healthcare, commerce, and business, recommender systems, museums and tourist scenarios, and group decision making. Among technologies, we focus on ambient data management and artificial intelligence; for example planning, learning, event-condition-action rules, temporal reasoning, and agent-oriented technologies. The survey is not intended to be exhaustive, but to convey a broad range of applications, technologies, and technical, social, and ethical challenges.

Categories and Subject Descriptors: I.2.0 [**Artificial Intelligence**]: General; I.2.1 [**Applications and Expert Systems**]; J.0 [**Computer Applications**]: General; A.1 [**General Literature**]: Introductory and Survey

General Terms: Design, Experimentation, Human Factors

Additional Key Words and Phrases: Ambient intelligence, assisted living, agents, multiagent systems, social and ethical issues

ACM Reference Format:
Sadri, F. 2011. Ambient intelligence: A survey. ACM Comput. Surv. 43, 4, Article 36 (October 2011), 66 pages.
DOI = 10.1145/1978802.1978815 http://doi.acm.org/10.1145/1978802.1978815

1. INTRODUCTION

The term Ambient Intelligence (AmI) was coined by the European Commission, when in 2001 one of its Programme Advisory Groups, the European Community's Information Society Technology (ISTAG), launched the AmI challenge (Ducatel et al. 2001), later updated in 2003 (Ducatel et al. 2003). But although the term AmI originated in Europe, its usage and intended goals have been adopted worldwide, with many related projects and research programs in recent years. This very active area has been the subject of several collections of papers, and special issue publications (e.g. Werner et al. [2005], Cai and Abascal [2006], Augusto and Nugent [2006], Aarts et al. [2009], IEEE Intelligent Systems Magazine [2008]) and several specialized workshops and conferences. Some examples are: the European Conference on Ambient Intelligence which has been running for the past three years, the International Conference on Ubiquitous Robots and Ambient Intelligence, running for the past 6 years, the Ambient Intelligence Forum 2009, and the International Scientific Conferences in the Ambience series held in 2005 and 2008.

Ambient Intelligence is the vision of a future in which environments support the people inhabiting them. This envisaged environment is unobtrusive, interconnected, adaptable, dynamic, embedded, and intelligent. In this vision the traditional computer input and output media disappear. Instead processors and sensors are integrated in

Author's address: Department of Computing, Imperial College London, 180 Queen's Gate, London SW7 2BZ, U.K.; email: fs@doc.ic.ac.uk.
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2011 ACM 0360-0300/2011/10-ART36 $10.00
DOI 10.1145/1978802.1978815 http://doi.acm.org/10.1145/1978802.1978815

everyday objects. So for example, instead of using mice, screens, and keyboards, we may communicate directly with our clothes, household devices, and furniture, and these may communicate with each other and with other people's devices and furniture. The envisioned AmI environment is sensitive to the needs of its inhabitants, and capable of anticipating their needs and behavior. It is aware of their personal requirements and preferences, and interacts with people in a user-friendly way, even capable of expressing, recognizing and responding to emotion.

Most authors broadly share similar views of the features required for AmI applications. Gaggioli [2005] identifies AmI as an intelligent, embedded, digital environment that is sensitive and responsive to the presence of people. The key features here are intelligence and embedding. By intelligence, he means the system is sensitive to context, is adaptive, learns from the behavior of users, and eventually, recognizes and expresses emotion. By embedding he means small, possibly minituarized devices that merge into the background of people's activities and environments. Similarly, Aarts [2004] identifies five related key technology features: embedded, context aware, personalized, adaptive, and anticipatory.

The AmI environment is based on miniaturized and low-cost hardware, providing complex networks of heterogeneous information appliances or smart artifacts. These, individually or as an ensemble, can assist users in day-to-day or exceptional activities. They can demonstrate common sense and problem-solving ability, and assume responsibility for certain need, of the user, such as monitoring, alerting and helping users in their tasks. Advances in AmI are facilitated by the long-term increases in the power of microprocessors and nanotechnology, as well as the cost-efficiency of storage capacities and communication bandwidths.

Some of the more modest visions of AmI are near to realization, and others look more like science fiction for now. But the scope of potential applications is vast, and in this survey we look at many. AmI has the potential of freeing people from tedious regular routine tasks. It can provide assistance in many circumstances. For example parents may never loose track of their children in crowds, because of location sensors and miniature communication devices sewn into the fabric of clothes. Blind people may be guided in unfamiliar environments by intelligent signposts and public transport timetables that may communicate via wireless headsets [Coroama and Rothenbacher 2003]. Our washing machines may query our dirty clothes for the required washing programs. Traditional memory aids can remind the user about activities on their daily schedule, but more sophisticated memory aids, on the other hand, can be context-sensitive. They can observe the user in their activities, guess their desired tasks and on that basis issue reminders and guidance [Philipose et al. 2004, Pollock et al. 2003].

AmI has potential applications in many areas of life, including in the home, office, transport, and industry; entertainment, tourism, recommender systems, safety systems, ehealth, and supported living of many different variations.

The potential of AmI at home has been the subject of research for at least a decade in some major industries. The Microsoft Corporation Domestic Home in Redmond, Washington, for example, showcases AmI-based smart appliance technologies for the home. There are no desktop or laptop computers, but the wallpaper is interactive, and can be controlled by tablet PCs. The mailbox outside tracks the mailman's location using GPS, and users can get a real-time estimate of when mail will arrive, on the mailbox display or by cell phone. RFID (radio frequency identification) tags embedded into envelopes even provide information about what mail is on the way. RFID tags are active barcodes, which attach digital information to objects. This technology is being increasingly used by industry for tracking inventory. These tags can be quite small and do not require a battery. In the Microsoft Home there are RFID tags on clothes as well. In a bedroom, by holding clothes up to a mirror which doubles as a screen one can get information about them, including whether matching items like a skirt or jacket are in the wardrobe or the laundry. The kitchen has an intelligent note board. If you pin a pizza coupon on it the restaurant's menu and phone number are displayed. You can call the restaurant with a tap on the board.

In the UK, British Telecom (BT) and Liverpool City Council have run trails on telecare technology within a project devised by BT's Pervasive ICT Research Centre [BJHC&IM 2006]. The trial concerned a system that responded to crises in the home. Each home contained 12 wireless sensors, including passive infrared sensors, entry switch sensors, occupancy sensors, a toilet flush sensor, and a temperature sensor, all connected to a central gateway, and to a BT server, via a secure broadband IP channel. If a cause for concern is flagged, a voice call is made to the home's occupier. If he/she confirms that they are okay, voice recognition technology is used to cancel the concern, otherwise a voice alert is sent to selected personnel, who can then use a standard Web browser to access information about the inhabitant and the circumstances of the alert.

Personalization and context-awareness have been some of the objectives of AmI technologies developed at Philips Research laboratories.[1] In conjunction with sportswear manufacturer Nike. Philips has investigated biosensors that can be embedded in clothing to detect parameters such as heart rate, respiration rate, and blood oxygen levels. The aim is to provide personal fitness training and healthcare-monitoring within sports clothing. Another project is Philips Research's CAMP (Context Aware Messaging Platform), aimed at providing context awareness in mobile devices. Location is taken as an indicator of context, and a prototype system uses Bluetooth short-range wireless transmissions between stationary beacons and mobile devices. Beacons and mobile devices can transmit information, for example about their interests and offerings. Thus in a meeting with unfamiliar people, a beacon in the room could transmit their electronic business cards to your mobile device to help you identify them. Similarly, in a shopping center, a beacon could alert you via your mobile device about any friends that are nearby, and which shops they are in, as well as any special offers on products you are interested in.

Context awareness is also one of the key elements of the vision of ihospital proposed by Sanchez et al. [2008]. The ihopspital, an interactive, smart hospital environment, provides context-aware communication, whereby for example, a message can be sent to the doctor responsible for a patient in the next shift, without the need to know who the doctor will be. Information about context, such as location, time, and the roles of the people present, and RFID-tagged artifacts, is also used for activity recognition; for example the activities of a ward nurse and doctor during the morning hours, handling reports and case files, may be interpreted as patient care or clinical case-assessment.

The use of AmI for leisure and tourism has been explored by Augello et al. [2007], Gonzales et al. [2004b], Park et al. [2006] and Rumetshofer et al. [2003], amongst others. It has been explored also in several EU-funded projects such as AmbieSense (2002-2005), and FP6 e-Sense (2006-2008, http://www.ist-e-sense.org/), and FP7 Sensei (2008-2011, http://www.sensei-project.eu/). Busetta et al. [2003, 2004], within projects PEACH and TICCA, have explored the use of AmI for interactive museums. In their scenarios museum exhibits are equipped with several agent-controlled devices that are capable of producing a presentation about the exhibit but with different capabilities, e.g. pictures, video, text, audio. The devices sense approaching visitors and produce presentations, controlled by a multi-agent architecture and role-based real-time coordination mechanism.

Research into ambient recommender systems, for example, Gonzales et al. [2004a, 2004b, 2005, 2006], attempts to extend the functionality of conventional recommender

[1] http://www.research.philips.com/technologies/syst_softw/ami/breakthroughs.html.

systems. This involves designing proactive systems that are intelligent and ubiquitous advisors in everyday contexts, which reduce the information overload for the user, and are sensitive, adaptive and responsive to user needs, habits and emotions. Recognizing human emotions and using this recognition to reduce user frustration, for example in games or tutoring systems, are also active areas of research within AmI (e.g. Picard 2000 and 2007, and Paiva et al. 2007). Recognizing and utilizing emotions have also been considered in other areas such as group decision making [Marreiros et al. 2007a, 2007b].

Inevitably the social impact and the acceptance of such potentially powerful and intrusive technologies are topics of debate, and have been since their conception [Weiser 1993]. The success and acceptance of AmI by the public will depend on how secure and reliable it is and to what extent it is perceived to allow the protection of the rights and privacy of individuals.

As well as setting out a technological vision of AmI, the European Commission (ISTAG) report [Ducatel et al., 2001] identified several characteristics necessary for its social acceptance. For example it suggested that AmI systems should:

—facilitate human contact;
—be oriented towards community and cultural enhancement;
—inspire trust and confidence;
—be controllable by ordinary people—there should be an off-switch within reach.

This survey includes reports of some work carried out to explore such issues, including the impact of AmI on social norms, as well as its possible economic implications.

The AmI vision may be thought of as the convergence of at least three areas of computing: ubiquitous computing, sensor network technology, and artificial intelligence. Mark Weiser [1991] describes ubiquitous computing as the concept of computers weaving "themselves into the fabric of everyday life until they are undistinguishable from it." He believes that ubiquity is the key to providing effective, flexible, and convenient ambient environments, and outlines scenarios to this effect that use a multiple of devices combined with simple context-awareness, provided by identification tags for people and objects. But AmI is now a very broad multidisciplinary endeavor, drawing on, and consequently enhancing, existing technologies, including ubiquitous computing, ubiquitous communication, intelligent user-friendly interfaces (e.g. multimodal, visual, sound, and speech), artificial intelligence and multiagent systems. It is further supported by sensor network technologies, location, and motion trackers, medical devices, decision support systems, mobile communications, and wireless networks.

Progress in wireless technologies, sensor networks, display capabilities, processing speeds, and mobile services has paved the way to impressive developments in the AmI field. These developments help provide much useful raw information for AmI applications, and artificial intelligence and agent-based technologies have been explored to take full advantage of such information in order to provide the degree of intelligence, flexibility, and naturalness envisaged. Many AI techniques have been explored in this context, for example learning (e.g. Cook et al. [2006]), affective computing (e.g. Picard [2000]), case-based reasoning (e.g. Kofod-Peterson and Aamodt [2006]), temporal reasoning (e.g. Augusto et al. [2005]), planning (e.g. Amigone et al. [2005]), decision trees (e.g. Stankovski and Trnkoczy [2006]), and fuzzy logics (e.g. Hagras et al. [2004]). Some database techniques have also been used, such as event-condition-action rules (Augusto et al. [2004, 2005, 2008]), and (extensions of) SQL-based data management techniques (e.g. Feng et al. [2004]).

An emergent field is the multidisciplinary synergy of ambient intelligence and natural language processing (NLP), sometimes called Ambient Semantic Computing (see, for example, IJSC [2009]). This field refers to the application of NLP, ontology-related work and computer perception techniques such as situation detection and speech recognition, to ambient intelligence, to combine naturalness with the smartness of AmI. In a roadmap article, Gurevych and Mühlhäuser [2007] highlight the mutual benefits that NLP and AmI can bring to each other. On the one hand, NLP has the potential of transforming human-oriented unstructured information into structured knowledge, and of providing natural language and speech interfaces. On the other hand, AmI techniques can provide sensed contextual information for disambiguating natural language.

Agent technology is commonly used in AmI applications. This is not surprising, as the pervasive nature of AmI requires distributed information and problem solving, and agents are known to facilitate such architectures. Agents can be used as useful abstractions in AmI systems, for example for devices and functionalities, and as paradigms for implementation. They can be used at various levels: for example, to model individual devices present in the AmI environment, or they can be exploited at the middleware level to coordinate the activities of the lower level entities, or they can be used at a higher level to form the interface for humans.

Favela et al. [2004] and Rodriguez et al. [2004, 2005], for example, describe an architecture called SALSA, for health care, which uses agents as abstractions, to act on behalf of users, to represent services, and to provide wrapping of complex functionality to be hidden from the user. Amigone et al. [2005] describe a mixed central and distributed approach to planning in an AmI environment where devices are represented by agents, and enter and leave the environment. Da Silva et al. [2007] describe an architecture for an AmI-enhanced bookshop, where there is a one-to-one map from the elements of the bookshop (e.g. customers and audio-visual devices) onto agents in a parallel software world of agents. Misker et al. [2004] describe a prototype where devices and services in an ambient intelligence setting are represented by agents. The aim of the prototype is to experiment with issues in user interaction with agents, with particular attention to the balance between the autonomy of agents and the requirement for the user to be in charge.

Different architectures have been used for organizing agents in AmI systems. For example, Da Silva et al. [2007] describe a blackboard architecture whereby administrative agents look for messages and then contact appropriate service agents to deal with each message. LoudVoice [Busetta et al. 2004], on the other hand, is based on an architecture of implicit organizations emerging through role-based communication. Several other AmI proposals [e.g. Augusto et al. 2004, 2005, 2008] use a centralized agent architecture, where a single agent receives all the information and is responsible for all the decision-making. Robocare [Cesta et al. 2002, 2005], for example, has an event-manager agent that processes all requests and then directs each request to an appropriate agent via agent-to-agent communication.

Other non-agent-based generic architectures have also been proposed. For example, Piva et al. [2005] and Marchesotti et al. [2005] propose architectures inspired by a model of human consciousness given in Damasio [2000]. Ramos [2008] also proposes a non-agent-based architecture for ambient intelligence systems that combines the use of AI techniques, such as planning, learning, and modelling, with operational parts such as sensing and actuating. He gives brief example applications in group decision-making, assistance of pilots in flights, and restoration of power after incidents of power failure.

This survey does not attempt to be exhaustive in capturing all possible disciplines and reported work within the AmI framework. In particular, it does not address the technologies of sensor networks and hardware issues. It focuses primarily on applications and the intelligence aspects of AmI. It is biased towards the AI and agent technologies used in AmI, although where attempting to give a broader flavor of the field of work it touches on other prominent approaches and technologies.

This article is structured as follows. First, in Sections 2–6, we explore several major applications of AmI. For each broad application we review a collection of papers, and where necessary for understanding a technical contribution, we provide brief background material on the technology used, for example the BDI architecture, abductive logic programming, and event-condition-action rules. Then, in Section 7, we look in detail at additional data management and AI technologies in development for use in AmI systems. In Section 8 we look at affective computing and the consideration of human emotions in ambient intelligence applications, and in Section 9 we consider the social and ethical issues. Finally, in section 10 we conclude the article.

2. AMI APPLICATION: AMI AT HOME

A great deal of the vision of AmI is targeted at the home environment. AmI at home can provide context-awareness and proactiveness to support everyday living. For example [Encarnacao and Kirste 2005] the bathroom mirror can remind a person of the medication he/she has to take, and the car stereo can tune into the same station that was on during breakfast. Some ideas and functionalities are distant visions, such as self-monitoring and self-painting walls, and lighting and furniture recognizing emotions and moods; some have already produced prototypes, such as dormitories that learn simple preferences of their single occupiers regarding open or shut windows and level of lighting or heating. Some simple AmI-based devices are even routinely commercially available, such as temperature sensitive heating systems, movement-sensitive lighting, and light-sensitive blinds.

Various names have been used to describe homes equipped with pervasive technology to provide AmI services to the inhabitants. *Smart Homes* may be the most popular term, and others terms include *aware houses*, *intelligent homes*, *integrated environments*, *alive, interactive, responsive homes/environments*. Innovation in domestic technology has long been driven and marketed by the desire to reduce labor and improve the quality of time spent at home. This continues to be one of the motivations for development of AmI at home. Other factors include technological advances and expectations, and an increasing trend in a way of life that blurs the boundaries between home, work, and places of rest and entertainment.

2.1. What are Smart Homes?

A Smart Home environment is a home equipped with sensors and activators of various types to monitor activities and movement, and to monitor risk situations, such as fire and smoke alarms. In general, there are usually three main components associated with Smart Homes, a set of sensors, a set of activators for controlling the sensors and other equipment, such as cookers, windows, and so on, and computing facilities to which the sensors and activators are linked.

The sensors can, for example, detect if a Fancet is on or not, a weight is on a seat, or whether someone wearing an identifying tag has passed under a door joining two rooms. The sensor information may require some form of further processing to produce useful data, for example about location (e.g. Rodriguez et al. [2004]—see Section 4.1 in this survey), or about the task the resident is trying to achieve (e.g. Geib and Goldman [2001], Geib [2002]). It can be used to trigger some form of intervention on the part of the Smart Home to avert disaster (e.g. Augusto et al. [2004, 2005, 2008]—see Section 3.4 in this survey). It can also be used to learn behavior patterns of the occupant (e.g. Hagras et al. [2004]—see Section 2.5 in this survey).

The Smart Home concept is often used to support people with some cognitive impairment living on their own. Such homes are intended to provide better quality of life and greater levels of independence, and to reduce the need for institutionalization by extending the time that people can live in their own homes.

Alarm systems are standard conventional equipment provided currently in most homes. But the Smart Home concept goes much beyond these, to envisage systems that can monitor the resident's activities, compare them with his/her profile, take action if necessary, and provide help and guidance regarding safety, health, and medication, as well as entertainment and other matters. For example, currently there are alarm buttons in the home or around a person's neck or wrist to send alert in the event of emergency, fire, unconsciousness, or a fall. But the problem is that there may not be enough time for the user to trigger the alarm [Abascal et al. 2008]. To alleviate this problem there is a move towards alarms that can activate themselves, for example based on vital functions such as pulse rates and blood pressure. Similarly we have many standard automatic systems that switch the light on when they detect motion, but a Smart Home can turn the light on in anticipation of its resident's needs, or even turn the light on as a reminder or cue to where the resident should be moving.

AmI services envisaged in the home environment include the following.

—performing many everyday tasks automatically, thus reducing the burden of managing the house, for example, controlling household appliances and other objects to make household duties and maintenance tasks easier;
—improving economy of usage of utilities, such as electricity, by controlling the lights and window blinds, for example;
—improving safety and security, for example by preventing accidents, recognizing and rapid reporting of accidents, tracking people and providing entry access control with sophisticated interfaces. Safety/security can include the obvious, such as entry access control and alarm facilities, automatic safety protection for appliances such as irons and ovens, to safety in terms of health and biomedical monitoring, to buildings that monitor themselves and alert about, or even make, necessary reconstructions and repairs;
—improving quality of life, for example through entertainment and increasing comfort levels;
—supporting independent living for the elderly and the impaired. This is addressed in the next section, as it is an important application in its own right.

The technologies needed are diverse. Two main general components are a networked set of sensors and computing facilities, to integrate and interpret the sensor information, and learn or initiate action. The sensors can be specialized for sensing carbon monoxide, heat, motion, or for detecting whether a door or window is open or shut. Some specific applications of AmI at home may require multidisciplinary collaboration, in computer science, electrical engineering, and possibly aspects of medicine and general health care, social sciences, and occupational therapy.

Much of the work on Smart Homes concentrates on the hardware, sensors, and devices, but several authors [Augusto and Nugent 2006a, Cesta et al. 2005] have argued that AI techniques can help the evolution of Smart Homes by bringing a degree of sophistication to the processing of information provided by the devices and sensors. This processing can be focused on learning user profiles (for example Cook et al. [2006]), on diagnostic capabilities for determining whether a deviation from routine is a matter of concern, on advising whether or not a rescheduling of normal activities is possible after a deviation (for example Bahadori et al. [2004c]), and on combining elements of temporal and probabilistic reasoning to provide a more powerful setting for monitoring and intervening (for example Augusto and Nugent 2006a). Others (for example Sadri and Stathis [2008]) have argued that a rational and reactive agent architecture may be appropriated to use for smart home applications, as such an architecture is designed for dynamic environments and can be used to provide the required reasoning and adaptability.

It has also been argued (for example by Augusto and Nugent [2006a] and Patkos et al. [2007]) that smart home applications and AmI, in general, can contribute to the advancement of AI technologies, by providing AI with nontrivial testbeds and scenarios, requiring integration and development of several techniques.

We start the survey of AmI at home with an overview scene-setting paper, Friedewald et al. [2005], and then move on to more technical papers. The papers describe three implemented home systems, GENIO, MavHome, and iDorm, illustrating advances in home ambient systems and the use of RFID tags, speech recognition, learning and fuzzy logics in this application.

2.2. Friedewald et al. [2005]

This article discusses the opportunities that ambient intelligence can provide in the home environment, and describes the risks, and the balance between the technology enhancing or overburdening life. In their view AmI solutions at home will cover four areas.

(a) Home automation
(b) Communication and socialization
(c) Rest, refreshing, and sport
(d) Household work and learning.

Home automation. This includes the control of electricity, heating, air-conditioning, ventilation, and fire and intruder alarms. Many such functions already exist, for example, motion detectors, sensitive air conditioners, automatic heating adjusters, and automatic sunlight-sensitive blinds. Several challenges have been identified for enhancing such functionalities in intelligent homes. One is in providing more flexible interfaces for such functions, such as voice, hand gesture, and facial expression. Another is the adjustment of such functions according to the home's knowledge of the preferences of its residents. A third is in recognizing the user's likes or dislikes on the fly from their voices or facial expressions, and a fourth is in managing spaces shared by several occupiers, for example, friends and families. Maintaining security at home can go beyond ordinary access control to locks capable of identifying persons, hands-free unlocking, and affordable face/voice/iris recognition facilities. Such recognition techniques could also be used in other aspects of the home, for example to provide the preferred entertainment for the person occupying a room.

Communication and socialization. We already enjoy technological advances in fast access and communication via the Internet, with hand-held or hands-free devices, and with mobile photo and video facilities. AmI technological advances could provide access to digitalized documents, family photos, and films, regardless of location and equipment. They could also open up different ways of participating in the civil society, for example by e-participation and electronic voting in politics, referenda (e-governments), and unions. Furthermore, and possibly in the longer term, AmI may allow formation of ad hoc wireless social contacts and communities of people with shared interests and needs through spontaneous contacts of their virtual profiles, for example, people wanting to commute to a particular area.

Rest, refreshing, and sport. Resting can be supported by AmI through multifunctional, flexible furniture, with sensors embedded in the furniture measuring blood pressure and pulse for example, together with bio-identifiers recognizing the occupant of the furniture and recognizing their preferences, for example for a nap, massage, soft music, and so on. Databases can recognize and play a piece of music after the occupant hums a few bars of it. The home may contain AmI facilities dedicated to health and fitness, for example counting the amount of time spent using the stairs and estimating the calories used, as well as biomedical monitors to regularly check blood pressure and heart rates.

Household work and learning. Household appliances already reduce the amount of housework. AmI-embedded appliances can go further. For example washing machines can determine the program required for washing by checking the degree of soil of the clothes, vacuum cleaning robots can not only navigate around obstructions, but can also recognize small items and tell apart, for example, a small bottle top and an ear-ring on the floor. It could even be possible to discard cleaning appliances and embed cleaning properties within materials themselves. The refrigerator can become a kitchen computer displaying information about its contents, and listing missing items. The home AmI vision extends to entire buildings, where buildings have self-repairing elements: walls, floors, and ceilings monitoring the state of repair and warning when there is need for work, walls self-painting, gardens recognizing the need for mowing, and lawn mowing robots recognizing flowers that are to be preserved in their path.

Clearly realizing such visions presents major challenges, among them the context sensitive and personalized requirements of AmI (how to accurately recognize what the human inhabitant needs or wants). There are also social and cognitive challenges, for example how to provide the functionalities in such a way that they enhance life rather than overburden it, and how to ensure that they do not limit initiative and weaken the cognitive abilities of more vulnerable people, such as the elderly.

2.3. Garate et al. [2005] GENIO

This paper describes the GENIO project, which is a collaborative effort between two companies in Spain, Fagor Electrodomesticos and Ikerlan, aimed at providing ambient intelligence at home. GENIO is Spanish for Genie, the one that grants your desires. In this work household appliances such as fridge, washing machine, oven, sensors, security, and heating devices are networked and managed by a central controller which can respond to, and hold, a dialogue with the user in natural language (Spanish). This central domotic controller is called Maior-Domo and has a visible avatar representation. A demonstrator has been built consisting of a kitchen and sitting room with various domestic and entertainment appliances controlled by Maior-Domo. Some of the appliances and devices and their capabilities are as follows.

—*Oven.* This has a database of recipes and once a recipe is chosen it can control the oven temperature, timing and, method of heating.
—*Fridge.* This has an RFID (Radio Frequency Identification) antenna and an RFID reader inside, allowing it to read the goods stored in it. The goods have to be equipped with RFID tags.
—A large panel containing the electronic circuit boards of a washing machine and a dishwasher. These circuit boards substitute for the actual appliances for the purposes of the demonstrator.
—*A computer working as Maior-Domo.* This can communicate with the oven, washing machine, and dishwasher electronically. It is also equipped with a speech recognition system (developed by means of Java, VXML, JSP, JavaBeans), a Text to Speech component, applications to command the appliances, and digital information such as photos, songs, and videos.
—A pocket microphone for the user, allowing wireless connection to Maior-Domo from anywhere in the environment.

Some example scenarios used in the demonstrator are the following.

—*Reading emails.* User: “Maior-Domo, how many emails do I have.” Maior-Domo answers with a number and gives the sender of each one, from which the user can choose which one Maior-Domo should read for him.

—*Activating washing machine.* User asks Maior-Domo to wash his clothes by the time he gets back from work. Maior-Domo inquires when he will be back from work and which washing program is required, and on getting the answers sets the washing machine.
—*Checking goods in the fridge and preparing a shopping list.* User specifies a list of items he would always like to purchase, and can command Maior-Domo, by voice, to add further items to the list and download it.
—*Preparing a recipe.* User can ask Maior-Domo for a recipe, for roast chicken say, and specify the number of people it is required for. Maior-Domo adjusts the recipe to the number required, and can on command, read the recipe line by line, pausing for a user command after each line.
—*Entertainment.* User can ask for some Mozart music, for example. Maior-Domo lists what is available and plays what the user selects, and stops on command.

The major contributions of this work are the real settings, the demonstrator allowing control of several appliances, and the voice processing. The proposed future work includes recognition of the person giving commands to Maior-Domo, for example to ensure that a child does not activate the security system or the oven.

2.4. Cook et al. [2006] MavHome

This paper describes the MavHome (Managing An Intelligent Versatile Home) project at the University of Texas at Arlington. The objective of the project is to "create a home that acts as a rational agent," that has sensors and effectors, and that acquires and applies information about the inhabitants to provide comfort and efficiency. Comfort is in terms of the ambience of the environment, including temperature, ventilation, and lighting. Efficiency is in terms of the cost of utilities, such as gas and electricity. Thus the aims are twofold, to reduce the need for inhabitants to intervene to make changes for themselves in the environment, and to reduce energy consumption.

A Mavhome idealistic scenario is as follows. In Bob's house:

> "At 6:45 am MavHome turns up the heat because it has learned that the home needs 15 minutes to warm to optimal temperature for waking. The alarm goes off at 7:00, which signals the bedroom light to go on as well as the coffee maker in the kitchen. Bob steps into the bathroom and turns on the light. Mavhome records this interaction, displays the morning news on the bathroom video screen, and turns on the shower. While Bob is shaving MavHome senses that Bob has gained two pounds over the last week. MavHome informs Bob that this has been a trend over the last two months and offers suggestions for changing his lifestyle. When Bob finishes grooming, the bathroom light turns off while the blinds in the kitchen and living room open (an energy saving alternative to Bob's normal approach of turning on the living room, kitchen and hallway lights). When Bob leaves for work, MavHome secures the home, lowers the temperature, starts the robot vaccum, and turns on the lawn sprinklers ...... MavHome tracks Bob's activities while he is away from home in order to inform him of problems at home and to have the house temperature and hot tub prepared for his return at 6:00."

The MavHome architecture is a hierarchy of rational cooperating agents, each agent consisting of 4 layers. These layers, from top to bottom, are: the Decision layer, which decides what action to execute based on information provided by other layers, the Information layer, which collects, stores, and generates knowledge for decision-making, the Communication layer, which is used for information-passing between agents, and between the house and external devices, and the Physical layer, which contains basic hardware, devices, and network within the house.

Information about the current state is passed on to the Decision layer via the Physical layer through the Communication and Information layers. Decisions about what actions have to be executed are passed on from the Decision layer to the Physical layer via the Information and Communication layers. In more detail, sensors collect data about the environment, and transmit the data to the agents through the communication layer. The data is stored in a database and new data can be passed on to the decision layer, which may decide that an action should be executed. The decision is passed on to the information layer and stored in the database and passed on to the appropriate effectors in the Physical layer for action execution.

Communication is both point-to-point and publish-subscribe, and uses the CORBA (Common Object Request Broker Architecture) model, which is an industry standard specification developed by the Object Management Group (OMG) to aid the creation and usage of distributed objects. All agents register their presence using Zero Configuration (zeroconf) technologies. Zeroconf provides for automatic configuration and address allocation on wireless networks, allowing agents to join dynamically.

Mavhome combines several machine learning algorithms to learn the inhabitant's habits (it is assumed the home is inhabited by one person at a time), predict their next action and form action policies for manipulating the environment. To learn the inhabitant's habits, data mining is applied to data collected from observations of their activities and interactions with the environment, to extract patterns. The data mining is based on an approach by Agrawal and Srikant [1995] that deals with time-ordered transactions. To predict the inhabitant's actions, a compression algorithm called LZ78 [Ziv and Lampel 1978] is used in conjunction with a history of the inhabitant's interactions with the devices in the environment. To learn action policies, reinforcement learning is applied with a model of the system as a Markov Decision Process (MDP), where states are associated with rewards and state transitions are probabilistic. Negative reinforcement is received when the inhabitant immediately reverses the automatic action of the system (e.g. turns the light back on).

All MavHome components are implemented and are being tested in two environments, a workplace environment and an on-campus apartment with a full time student occupant. In both environments the sensory information includes light, temperature, humidity, motion (via passive infrared sensors), seat (occupied or not), and door and window (open or closed) status. The apartment also has sensors for leaks, smoke detection, vent position, and CO detection. The aim of the experiments in both settings has been to minimise the inhabitant's manual interaction with devices. Variations in the learning algorithms are mapped against reductions in manual interactions. The results reported indicate a reduction of between 72% and 76% in both environments.

2.5. Hagras et al. [2004] iDorm

This work reports the use of fuzzy logic-based techniques to learn user preferences in an ordinary living environment. They have devised an experimental intelligent inhabited environment, called the iDorm (intelligent dormitory) at the University of Essex, UK. The iDorm contains space for various activities such as sleeping, working, and entertaining, and contains various items of furniture such as bed, desk, wardrobe, and multimedia entertainment system. It is fitted with multiple sensors and effectors. The sensors can sense temperature, occupancy (for example user sitting at desk, user lying in bed), humidity, and light levels. The effectors can open and close doors, and adjust heaters and blinds.

The iDorm has embedded computational components including:

(1) *iDorm embedded agent.* This receives readings from the sensors about time of day, room light and temperature, outside light and temperature, state of window (closed, open), and user's activity, for example sitting at desk, sitting, lying in bed, using entertainment system, computer, and so on. It contains the user's learned behavior, and on the basis of this and the sensor data, it computes any appropriate actions,

using a fuzzy logic technique. The actions include adjusting a fan heater/cooler, lights, or blinds.

(2) *A robot.* This is a physical robot under the control of the iDorm agent. The robot is equipped with some navigation capabilities, such as obstacle avoidance, and can move and carry items such as food, drink, and medicine. iDorm is aware of the robot's location and sends it instructions for moving to a destination and carrying objects.

iDorm deals with two types of rules, static (user-independent) rules such as how to react in an emergency, and to lower the lights and temperature when the room is unoccupied, and learned rules reflecting the user preferences. Each user is identified by a unique ID. When a new user enters the room there is a monitoring period to sense user activities. This provides examples for the learning phase. Learning is based on negative reinforcement, as it is assumed that users change something in their environment when they are dissatisfied with it. The learning uses a fuzzy logic-based technique called Incremental Synchronous Learning (ISL).

After monitoring and learning, the iDorm agent can take control of the environment. If the user behavior changes, the learning system may need to modify some of the rules, so it will go back to a learning phase in which there can be a repeated learning process. There is an Experience Bank, which stores the previous occupants' rules. After the initial monitoring phase the system tries to find a best match from the bank. So the learning starts from some initial rule base. The paper reports a limit of 450 on the number of stored rules.

The following is a report of the experimental results. A user occupied iDorm for 5.5 days (132 hours). The user could interact with the environment wirelessly via a portable device (HP IPAQ) providing a form of remote control interface. At the end of this period the number of rules learned over time was examined, the hypothesis being that if the system was successful in learning user preferences, then there would be fewer incidents of user intervention over time, and thus fewer rules would be learned over time. The study found that in the first 24 hours many rules were learned. Then the number of learned rules fell significantly until time 60. From time 60 to 72 there was an increase in the number of learned rules, and after time 72 there were no new learned rules. The increase during the period 60–72 was explained by the observation that the user was adding new activities to his behavior during this time. The authors conclude that the embedded agents significantly reduced the user's need to intervene, and the Experience Bank also reduced the time that ISL needed to learn user behavior. The future work includes multiuser environments and fully functioning apartments.

3. AMI APPLICATION: DOMESTIC CARE OF THE ELDERLY, ASSISTED LIVING

One application area of ambient intelligence that has attracted much attention is support for independent living for the elderly. Better health care and quality of life have caused an increase in longevity. In 2002 it was estimated [UN 2002] that in Europe the elderly made up 12–17% of the total population. The proportion of the UK population over 65 years old is now 15%, compared to 11% in 1951 and 5% in 1911.[2] In the UK this year the number of people over age 60 is greater than the number below age 16. Several studies have highlighted a widely anticipated continuing increase in the elderly population. In the US people over 65 years are the group of the population that is fastest growing in numbers, and by 2020 they are expected to represent 1 in 6 of the population [Corchado et al. 2008].

[2]Centre for Economic Policy Research, http://www.cepr.org/pubs/bulletin/meets/416.htm.

The following is an extract from a UN article [UN 2004].

> "Life expectancy is assumed to rise continuously, with no upper limit, though at a slowing pace dictated by recent country trends. By 2100, life expectancy is expected to vary across countries from 66 to 97 years, and by 2300 from 87 to 106 years. . . . . . population aging becomes a predominant demographic feature . . . . . . Assuming that retirement age worldwide in 2004 was 65 years, people retired on average only two weeks short of their life expectancy. Assuming that retirement age stays unchanged, by 2300 people will retire 31 years short of their life expectancy."

Another point made in the 2004 UN article is that "... the cost of supporting each aged person is greater than that of supporting a child, according to one calculation, a ratio of five to three. Much of this additional cost arises from higher health costs ... Additionally, many aged persons—most in the West—live in separate accommodations and not with their children. This arrangement is necessarily more expensive, especially when additional services, such as specialized health care, are added." The same article estimates that by 2050 the worldwide population of elderly people (over 60 years) will be larger than the population of children (0–14 years). The demographic changes are further illustrated by what the UN article calls the *potential support ratio* (PSR), the number of persons aged 15–64 years per one person aged 65 or older. Between 1950 and 2002 the PSR fell from 12 to 9, and is expected to fall to 4 by the middle of this century.

Such demographic changes provide challenges for healthcare and maintenance of quality of life. We have seen, and may continue to see, an increasing shift in resources from institutional care towards providing care at home and towards preventative care. It is hoped that AmI-supported home automation can provide much better quality of life for the elderly than institutions. It is also hoped that the preventative quality of such support can have further important advantages in terms of cost effectiveness, by reducing the need for providing dedicated carers and institutional services.

Thus the main aims of the AmI research in this area have been to support independent living for elderly people, primarily people who might be in the early stages of cognitive or physical impairment who may be able to live in their own homes and would prefer to do so, but may need monitoring for the sake of their safety and well-being. The following papers illustrate a broad range of activities in this application, from technology implemented systems, to more theoretical concepts, as well as studies of the attitudes of targeted user groups.

3.1. Bahadori et al. [2004a, 2004b], Cesta et al. [2005] RoboCare

Robocare is an ongoing project[3] with the long term goal of contributing to improving quality of life of elderly people living independently in their homes, through the use of AI technologies to realize cognitively enhanced embedded technology. "The RoboCare Domestic Envirionment (RDE) is an experimental setup which recreates a three-room flat. It is intended as a testbed environment in which to test the ability of the domotic technology built by the various research units, such as non-intrusive monitoring, domestic robots and environmental supervision."

The papers describe two contributions in this direction.

—A tracking system for people and robots, exploiting stereo vision; and
—a task execution-monitoring and supervision component.

The two components provide a prototype system deployed in the testbed domestic environment.

[3]http://robocare.istc.cnr.it.

The entire system is based on an e-service-oriented architecture, and is composed of several hardware and software agents, each providing a set of services. The two components are integrated via an e-service-oriented middleware, where all agents provide their functionalities as services. The activities of all agents are organized by an event manager, which processes all requests for services and directs them to the appropriate agents. The tracking and monitoring components together with the event manager provide the Active Supervision Framework (ASF).

The tracking system. The tracking system is the primary component to help recognize interesting situations and report them to the decision-making agents within the monitoring part. It uses stereo vision technology to determine the 3D position of people and to track them as they move. Also by extracting a set of features, such as height and size, from the images, it makes it possible to identify certain poses, for example if the tracked person is lying on the ground, or sitting in a chair.

This system, thought of as an agent, provides services including *whereis, whichpose, howlong-in-a-pose, howlong-here, howmany-persons, howmany-robots, robot-close-to-a-person, what-activity*. Robots are distinguished from people by equipping them with special tags. Moreover, to simplify the task, the authors exclude crowded environments, allowing no more than 2–3 persons, and monitor only a portion of the domestic environment. The stereo vision system (a pair of Firewire webcams) is placed in a fixed position and can thus make use of information about the background to distinguish people. The software agent, which processes the images, distinguishes between the foreground (people/robots) and the background, computes the 3D position of the foreground, and associates each 3D cluster of the points with a particular robot from its tag or a nonspecific person. A special algorithm is used to attempt to distinguish inanimate objects from still animate ones. So that for example, when a person sits at a table and puts a bottle on it, after a while the bottle, but not the person, becomes part of the background, because the edges of the images of the bottle show low activity.

The execution monitoring and supervision system. The monitoring agent has knowledge of the assisted person's usual schedule. This is a set of predefined activities with resource constraints and predefined durations and other temporal constraints, such as, one activity cannot start before another starts, or cannot start before the end of another, and so on. The example given includes 6 activities.

A1: breakfast, A2: lunch, A3:dinner, A4, A5, A6: taking medication, with temporal constraints specifying the earliest start times and latest end times for A1, A2, A3; minimum required lapsed time between the end of A1 and start of A2; and between the end of A2 and start of A3. Furthermore, A4 should be done within a predefined time lapse after A1, A5 should be done immediately before A2, and A6 should be done immediately after A3.

The execution monitoring and supervision system has two tasks: one is to recognize if the actual situation, as informed by the sensors and the tracking system, has diverged from the schedule, and if so, fire an appropriate rule for repairing this. Another task is to ensure that all scheduled activities are actually executed, and issue warnings or alarms if they are not.

The actual situation is compared with the expected schedule in terms of temporal and resource constraints, and delays, variations in durations, and resource breakdowns, are identified and dealt with. A delay is represented by inserting a new temporal relationship. A variation in an activity's duration is represented by replacing the activity with a new one with the new duration but with all other characteristics remaining the same. A resource breakdown is represented by inserting a new *ghost activity*, which eats up all the resources that are no longer available. The schedule thus augmented is then fed into the ISES procedure (Iterative Sampling Earliest Solution), which is a constraint-based method originally designed to solve Resource constraint project scheduling problems with time windows [Cesta et al. 2002]. If the temporal constraints are deemed to have become unsatisfiable, a warning is issued. If not, the resource constraints are checked, and if they have become conflicting a schedule revision is attempted.

In the example of, the A1–A6 activities, whether a person is eating is assumed to be determined by the tracking system recognizing the position of the person as seated at a table in a particular room. A delay in any of the activities A1–A3 will have an impact on the schedule of the other activities, and a substantial delay in any of the activities A1–A3 may mean that one or more of A4–A6 cannot be executed in times compatible with their temporal constraints.

The system also attempts to detect unexpected events happening and recognizing whether or not they have caused inconsistencies with the normal schedule. An exogenous event may inflict a temporal or resource inconsistency with the schedule, the former happens, for example, when the extra event forces some scheduled activity beyond its deadline; the latter happens when an activity is forced within a time frame when resources are not available for it.

3.2. Georgia Institute of Technology - Aware Home Research initiative (AHRI)[4]

The initiative has a three-floor, 5040-square-foot home as a residential laboratory for interdisciplinary research and evaluation. There are several objectives including designing interactive environments for the elderly and exploring the social implications of such technologies. Some examples of the work reported are as follows.

I.L.S.A. (The Independent LifeStyle Assistant) [Plocher and Kiff, 2003, Guralnik and Haigh 2002] passively monitors the behavior of inhabitants and alerts caregivers in cases of emergency, for example a fall. It is a multiagent system in a Jade [Bellifemine et al. 1999] environment. Each agent in the system is responsible for one aspect, for example monitoring use of medications, issuing reminders, monitoring mobility, coordinating responses, and learning patterns of behavior. There are also a phone agent and a database agent. The agents need to interact with one another to achieve their goals. For example, the medication agent generates a reminder for taking a medicine, this is processed by the response coordinator agent, which then sends it to the phone agent. Finally the phone agent delivers the reminder.

The agents have sensors and actuators, and may have plan libraries that they use for recognizing the intention of users from their activities and for choosing the system's response. A focus of the project has been on high level reasoning capabilities in three main respects, machine learning, goal recognition, and response generation. Machine learning is used to learn schedules of regular activities of the inhabitants, to build models of which sensor firings correspond to which activities, and to raise alerts when an activity occurs that is probabilistically unlikely. Goal recognition is used to deduce the goal of the inhabitant from their observed actions. This is based on a system called PHATT (Probabilistic Hostile Agent Task Tracker) [Geib and Goldman, 2001]. PHATT uses a plan library of simple hierarchical task network plans. One very simple one is shown in Figure 1 indicating that doing action a, then action b, and then action c, achieves goal S. The observations of the action executions provide an execution trace. This is used to form probabilistic hypotheses about the actor's goal, and these, in turn, generate sets of pending (expected) actions.

The response generation is done in three stages, first the individual domain agents generate context-free responses, that is responses based only on their own domain information. These are sent to the central response coordinating agent, which prioritizes

[4] http://www.cc.gatech.edu/fce/ahri/projects/index.html.

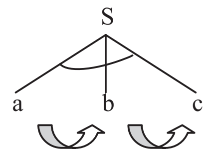

Fig. 1. A simple hierarchical task network.

them, thus injecting some degree of context-awareness, and also groups them. Finally, according to the decision of this agent, the response is directed to a device, for example the phone.

For field studies the I.L.S.A system was installed in several homes, including homes of system engineers and elderly clients. An evaluation of the system and the lessons learned, particularly in the use of agents and AI techniques, are given in Haigh and Kiff [2004] and Haigh et al. [2006]. The evaluation found that the use of agents did not provide the expected benefits. In particular, the need for coordination of the agents and centralized control outweighed the benefits of the distribution and independence of components provided by the agent architecture. The machine learning, on the other hand, was found to provide enhancements. In another paper [Geib 2002] complications that arise with goal recognition in the I.L.S.A. setting are discussed. These include recognizing multiple goals with interleaved execution of actions in their plans, and recognizing when goals are abandoned. Another complication is partial observability of actions. The latter two are particularly relevant in the elder care setting where plans may be abandoned due to forgetfulness and reminders need to be issued, and where the clients may not be in favor of having their every move observed.

Memory Mirror. This is deployed as a memory aid in the home and is intended to help with memory confusions that arise between the repeated episodes of frequent tasks, for example “Did I take my vitamin today or was that yesterday?”, “Has anyone fed the fish?”, “Did I take pain medication an hour ago, or did I decide to wait a bit longer?”.

The memory mirror provides a visual log of the movement and usage of specified objects during a period of time (e.g. 24 hours of a day). As a person uses an item, the usage is posted to the mirror and is recorded in a history log. The memory mirror also warns of possibly lost items as it tracks their removal and return from specific locations, normally their usual storage locations, such as medications on a nightstand. The memory mirror system uses RFID technology.

Gesture Pendant. This is a small pendant that the user can wear, which is equipped with a wireless camera. The system can monitor the user’s hand movements and can look for loss of motor skills and tremors, thus noticing onset of illness or problems with medication. It also provides hand movement interface to some devices, such as lighting and entertainment devices, reducing the need for dexterity and understanding of using switches and remote controls.

The Technology Coach. Elderly people often have to use home medical devices and the correct usage of such devices is crucial to their health. But these devices are often not designed with the cognitive issues that come with aging in mind, and the instructions are often created by marketing departments, rather than by designers. So older adults, typically, need initial training, and often follow-up daily assistance, to use these devices. The Technology Coach’s objective is to provide assistance in this, by providing ongoing feedback in using home medical devices. The system watches the use of the device via different tracking technologies and provides appropriate guidance. It combines two complimentary research efforts. One is aimed at understanding what kind of training and feedback should be used for older adults, by evaluating the use of conceptual and procedural feedback for both short-term and long-term use of the device [McLaughlin et al. 2002; Mykityshyn et al. 2002]. The other is to track and assess the use of the device by developing new techniques for modelling complex chronological tasks and new methods for recognizing actions from sensor data.

3.3. Niemela et al. [2007]

This paper reports a qualitative study of three ambient intelligence scenarios using focus groups of elderly people in two countries, Finland and Spain. The work is part of a project called MINAmI.[5] The project has the overall aims of developing and validating AmI applications that use mobile phones and adress the associated ethical and privacy issues. The applications concern health and safety monitoring at home, particularly for the elderly.

The three scenarios developed for the focus group study deal with monitoring the taking of medication, monitoring sleep, and home security. In these scenarios the mobile phone was used as the user's device for reading tags and sensors and for interacting with the AmI systems, the argument being that mobile phones are the best known and most commonly used mobile devices. The scenarios were animated as cartoons, each 1–3 minutes long. Two focus groups of elderly volunteers were used in Finland (one with 4 members and one with 10), and one in Spain (with 5 members). In addition, one focus group of medical experts was used in Finland and one in Spain. The elderly groups were mixed in sex and in their background experience with computers and mobile phones. The method of study was to allow a cartoon to be watched and then follow it by an open debate, where the participants were invited to take part in informal discussions of the presented material and any other related matters. Each scenario took 1–2 hours.

The first scenario involves a medication event monitoring system. In the cartoon Tom is given a smart pillbox by his doctor. The cap of the pillbox has a counter for opening and closing events and a clock. The pillbox communicates with a mobile phone by a wireless connection such as Bluetooth and can display on the mobile phone, the timed record of the cap openings and closings (assumed to represent the taking of medication). The data is also sent to an Internet database. If Tom forgets to take his pills for several days, the pillbox sends the information to a care center, which notifies the doctor.

The second scenario involves sleep monitoring, in particular sleep apnoea, a condition that causes interruptions in breathing during sleep. The monitoring is done via a sensor device that is worn on the forehead. The device detects EEG (electroencephalographic) brainwaves and movements of the head. The data can then be read, by taking a mobile phone near it, and sending the data to a healthcare center to be analyzed for abnormalities.

The third scenario deals with security at home. In the cartoon an elderly lady, Mrs S., lives in a big house equipped with sensors. When she moves to the stairs the lights come on automatically, and when the sensor system is set to security mode it can monitor movements in the house. For example when her dog goes to the basement a notification is sent to her mobile phone, and when a burglar attempts to enter the house through the basement an alarm is sounded in the house.

In the evaluation, the first scenario was judged the least favorable one. There were three reasons for it on the part of the elderly focus groups. They felt that it was too intrusive on their privacy, and that the data should not be reported to their doctor. They felt that relying on such devices would weaken people's cognitive abilities, and

[5] http://www.fp6-minami.org/.

they felt that the system would not be suitable for the majority of the elderly who take not just one medication but a variety. The expert focus groups were also concerned about privacy and data confidentiality, for example employers getting hold of the data and not funding private medical care for those employees who are deemed unreliable in taking their medication.

The second scenario was evaluated more positively by all groups, the elderly preferring the use of home equipment to having to wait in hospitals and sleeping in unfamiliar surroundings. The experts felt that the system could not replace hospital testing, but had some value in monitoring, and possibly cutting down on further hospital visits and use of resources, if it proved to be reliable. The third scenario was also viewed positively, where all involved felt it would lead to feeling safe, but all expressed concerns about the initial and maintenance costs. Overall, mobile phones were deemed suitable and feasible devices for user interfaces with AmI systems. It is interesting to note that another elder assisted living project, Roberta [Ballegaard et al. 2006], had based the user interface on tablet PCs, and in the evaluation found that few of the elders liked them or managed to learn to use them despite training.

The Niemela et al. [2007] study also looked at possible cultural differences in attitudes towards AmI. The Finnish groups were more concerned than the Spanish about privacy and security of the data collected via such systems. The Spanish, on the other hand, were more concerned than the Finnish about whether the home care and security centers could be trusted to access homes in alarm situations.

3.4. Event-Condition-Action (ECA) Rules for Smart Homes for elderly care

Several papers from one group of researchers [Augusto et al. 2008, 2005, 2004; Augusto and Nugent 2004, 2006a; Liu et al. 2006] make use of ECA rules and various extensions of them for applications in Smart Homes and supported living for the elderly.

ECA rules have the form:

On <event expression>
If <condition>
Do <action>.

The intuitive reading of such rules is that on detecting certain events, if certain conditions are true then certain actions should be executed. The event part (first line) is the trigger of the ECA rule. The rule is triggered if an event occurs that matches the event part of the rule. Then if the condition (second line) of a triggered rule is true the rule fires, requiring execution of the action (third line).

ECA rules were first proposed as part of Active Database technology [Paton and Diaz 1999]. In that context, events can be (single, disjunctions, or compositions of) database updates, conditions can be queries to be evaluated on the database, and actions can be (a single or a sequence of) database updates or calls to external programs or procedures. The execution of an action of one rule can trigger other rules. ECA rules are used in active databases for enforcement of integrity constraints, triggers and alerts, access constraints, gathering of statistics, and other applications.

In the paper forming the introduction of the book [Augusto and Nugent 2006a], the authors advocate the use of AI techniques to augment sensors and hardware-oriented technologies for Smart Home applications. They suggest that Smart Home applications can be enriched by AI techniques, and moreover that Smart Homes can provide a good range of applications and test beds for AI, which is typically less complex and less computationally prohibitive than many more conventional AI problems.

Here and in all the papers mentioned in the following, on the group's work, their proposed setting is a house or an apartment in a typical residential care institution, which provides independent but supported living. The support is provided by sensors and alarms connected to a central monitoring facility that performs all the reasoning. It is assumed that the person resident in the apartment wears a tag and that appliances have sensors and remote activators. There are sensors, for example, for heating, doorbell, phone, and so on; and medical facilities that can feed information into the central monitoring system about various statistics such as blood pressure and glucose level. The sensors are assumed to record, in effect, information about the movements of the resident, possibly with time stamps, and duration of stops at each location. The papers do not deal with sensor technology and data interpretation, but assume that the sensors provide information that can be fed into the ECA rules.

Augusto et al. [2005] advocate augmenting ECA rules with temporal features to facilitate the following two functionalities.

—Monitoring patients and reacting to observations; and
—longer term observations to help towards lifestyle profiling to be used by carers.

For the first functionality, the ECA rules are aimed at recognizing when a situation is hazardous, potentially so, or nonhazardous. An example of a hazardous situation is when the user (the elderly occupier) turns the cooker on, leaves the kitchen and does not return to the kitchen after a long lapse of time:

ON cooker has been turned on in the past and patient leaves kitchen
IF patient does not come back for a 'long' period of time
THEN turn cooker off and notify carers.

Augusto et al. [2004] also consider similar scenarios and ECA rules using a temporal language and reasoning engine that allows nonprimitive and durative events [Gomez et al. 2001], although the work does not address the issue of recognition of complex events. Some other examples are:

On blood pressure higher than 200/175 for more than two successive samples within the same day
If medication regimen to control blood pressure has been altered recently
Do notify clinical staff.

On person having been in bed for a long period of time
If not expected to be in bed during this period
Do contact relevant carer.

On leaving the house
If the person should be remaining in the house
Do raise alarm.

For the second functionality—longer term observations for lifestyle profiling, the papers do not provide technical details, but Augusto et al. [2005] outlines general possibilities. The idea here is to compile a care plan that includes information about the special needs of the inhabitant, their regular medication, dietary and other habits and routines, and reports of incidents, how they were dealt with and their outcomes. This care plan can then provide information for the designers of the ECA rules for monitoring and taking actions. This construction may not be automatic and no details are given. An example is that if the care plan indicates that the patient is expected to prepare her own food regularly, then an event of detecting that the fridge is not being used may indicate that she is skipping meals and may warrant a visit from a care worker. In general the informal care plan can be considered as an informal specification of the ECA rules. This can help identify the events the formal ECA system needs to focus on, the contexts (conditions) that may indicate risk, and the actions that may be appropriate. No implementation or evaluation details are given for either functionality.

Another paper [Augusto and Nugent 2004] also deals with ECA rules with temporal features for smart homes. It is assumed that several events can occur at the same time, and thus have the same time stamps, and several rules can fire at the same time, and if they do, all their actions will be executed. The actions may be external, giving instructions to carers, for example, or internal, causing changes to the knowledge base. Changes to the knowledge base (state) are dealt with by removing the old state and adding the new one. For example if an old state indicates "normal-context" and the triggering of an ECA rule indicates a "potentially-hazardous-context", this new information replaces the old. Persistence is assumed by default and negation is dealt with as negation-as-failure, that is unless there is a record of it in the database, the system assumes that the event or state has not happened. The temporal aspects are based on Galton [2004]. Two temporal operators are exemplified in the paper.

—ingr(S) indicates an event of ingression from a state ¬S to S. For the authors, state appears to mean a property. For example occurs(ingr(cooker-in-use), 1][2) is the notation that indicates that at time 1 the cooker was not on and at time 2 it was.
—trans(S1, S2) indicates an event of transition from a state S1 to an incompatible state S2. For example occurs(trans(at-kitchen, at-reception),2][3) denotes that there has been a sequence of sensor readings detecting that the person has moved from the kitchen to the reception at time 2.

The following is an example of an ECA rule using these operators and regarding the hazard of turning on the cooker and leaving the kitchen for an extended period of time.

On (occurs(ingr(cooker-in-use), T1][T2) and occurs(trans(at-kitchen, at-reception),T3][T4))
If (earlier(T2, T4) and ¬ holds(at-kitchen, [T2, Now]) and moreThanUnitsElapsed(T4, Now, 10 Mins))
Then (ApplyPossibleHazardProcedure and TryContact)

which has the intended reading that if there is an ingression to a state where the cooker is in use, followed by the person leaving the kitchen and not returning for more than 10 units of time then apply the "procedure to deal with potential hazard" (not defined) and also try contact.

Later papers [Liu et al. 2006, Augusto et al. 2008] advocate combining ECA rules with uncertainty as well as with spatio-temporal reasoning for smart home applications. For this purpose they use an existing technology called *rule-base inference methodology using the Evidential Reasoning* (RIMER) [Yang et al. 2006] for reasoning with uncertainty. They extend this with two temporal operators, *a ANDlater b* to mean a was detected before b, *a ANDsim b* to mean a and b were detected at the same time.

a *ANDlater* b is a shorthand for
∃ t1, t2 such that t1 < t2 and a is true at t1 and b is true at t2, and
a *ANDsim* b is a shorthand for
∃t such that a is true at t and b is true at t.

The authors argue that as the ECA rules have to deal with data collected from sensors they have to cater for uncertainty. Some sources of uncertainty are the following.

—Uncertainty in recognizing events (e.g. it is most likely that the patient has fallen asleep);
—Uncertainty in conditions (e.g. a sensor is considered to be activated with "high" confidence);
—Uncertainty in associating actions with events and conditions in ECA rules (e.g. if some activity is detected with "high" confidence, followed-by no movement detected for 10 units of time, then assume/record with 80% confidence that the patient has fainted).

An architecture is sketched whereby sensor and other data and the rule base are input into the RIMER engine, which attaches weights, priorities, and degrees of confidence to the components of the fired rules, and produces the resulting actions.

Augusto et al. [2008] elaborate on Liu et al. [2006] by giving further details on reasoning with temporal and incomplete information in the context of ECA rules for smart homes. They focus on an example of a smart home occupant being detected in one room and then under the doorway joining this room to another, and then no further movements are detected. The scenario requires reasoning with uncertain information, for example in inferring with varying degrees of confidence whether or not the occupant has fainted, depending on the degrees of confidence in the evidence of the sensor information. In this work two types of actions can form the action part of the ECA rules: recommendations to personnel, for example, *visit occupant*, and updates to the database of the system, for example, *assume occupant moved from reception to living room*. The time stamps of events are provided by regular sensor readings, each sensor reading provides a passage to the next time. To accommodate this, the system allows "no-event" inputs from sensor readings when the sensors detect no events. The knowledge base consists of ECA rules with confidence factors and weights. For example, a simplified version of a rule is the following.

IF at_kitchen_on with (H) *ANDlater*
under_door_on with (L) *ANDlater*
no_movement_detected with (H)
THEN estimated confidence that the occupant is compromised is
{(H,0), (M, 0.4), (L, 0.6), (N,0)}

In this ECA rule, the antecedents state that the sensors have indicated that the occupant (implicit here) is in the kitchen (at_kitchen_on), and there is high (H) confidence in this data, later he is detected under the door (*under_door_on*), and there is low (L) confidence in this data, and later no movement is detected with high confidence. The latter condition appears to be implemented as at_reception_off with H, which states that there is high confidence the occupant is not in the reception room. Presumably the door leads from the kitchen to the reception room. The conclusion represents that there is 40% confidence that there is medium (M) possibility that the occupant is compromised (e.g. fainted or fallen), 60% confidence that there is low (L) possibility, 0% confidence that there is high (H) possibility and 0% confidence that there is no (N) possibility.

The inputs to the system are given as, the followings, for example.

at_kitchen_on with {(H, 0.9), (M, 0.1), (L, 0), (N,0)} *ANDlater*
under_door_on with {(H, 0.9), (M, 0.1), (L, 0), (N,0)} *ANDlater*
at_reception_off with {(H, 1), (M, 0), (L, 0), (N,0)}.

An algorithm is described that first uses such inputs to calculate activation weights for all the rules with matching antecedents (in this example all those rules with *at_kitchen_on, under_door_on,* or *at_reception_off* in the antecedents). Intuitively, the closer the input is to the antecedents the higher the activation weight of the rule, where the closeness is dependent on the confidence factors. The activation weight of the given rule (assuming *no_movement_detected* is replaced by *at_reception_off*), for example is very low (in fact 0), given these inputs. Intuitively this is because of the mismatch of *under_door_on with (L)* in the antecedents and *under_door_on with* {*(H, 0.9), (M, 0.1), (L, 0), (N,0)*} in the inputs. Next the algorithm combines the conclusions of the activated rules according to the activation weights to compute a final overall conclusion, by a form of conflict resolution.

3.5. Abductive Logic Programming for Smart Homes for Elderly Care

Proposals have been made by Sadri [2007] and Sadri and Stathis [2008] for the use of abductive logic programming [Kakas et al. 1992] within an agent model for AmI applications for the elderly.

3.5.1. What is Abductive Logic Programming?. An abductive logic program (ALP) is a tuple <*P, IC, A*> where: *P* is a logic program, consisting of clauses of the form *Head if Body*, where *Head* is an atom and *Body* is a conjunction of literals (atoms or their negation). All variables in *Head* and *Body* are assumed universally quantified over the whole clause. *IC* is a set of integrity constraints of the form $L_1, \ldots, L_n \rightarrow H_1, \ldots, H_m$, where the *Li* are literals, the *Hi* are atoms, possibly *false*, and the "," on both sides of $\rightarrow$ denotes "*and*". All variables occurring only in the Hi are assumed existentially quantified on the right side of $\rightarrow$, and all others are assumed universally quantified over the integrity constraint. *A* is the set of abducible predicates. Abducible predicates do not occur in the heads of clauses, but can occur in the bodies and anywhere in the integrity constraints.

Given an ALP, *P*, and a query, *Q* (conjunction of literals signifying goals or observations), the purpose of abduction is to find a set of atoms in the abducible predicates (an abductive answer) that, in conjunction with *P*, entails *Q* while satisfying the *IC*, with respect to some notion of satisfaction and entailment. If Q is a goal then the abductive answer is a solution for achieving it, and if Q is an observation the abductive answer is a possible explanation for it.

The logic program in the ALP can be used to represent domain knowledge and plan libraries, and the integrity constraints can be used to model reactive rules, including ECA rules, and interaction policies—see, for example, Sadri et al. [2002] or Sadri and Toni [2000]. The use of integrity constraints as reactive rules provides active behavior similar to ECA rules, but with added declarative semantics, and with the background knowledge available via the logic program. The following are some examples of integrity constraints. The first two are reactive rules, and the third is an interaction policy.

unlocked-door(T), T $\geq$ 22:00 $\rightarrow$ tell(alert(unlocked(front-door)), T),
do(turn-on-light,T)

specifies that if the door is unlocked after 22:00 hours an alert should be given and the light should be turned on.

observed(on(gas), T), observed(¬lit-hob, T), observed(¬lit-hob, T + 5minutes)
$\rightarrow$ do(turn-off(gas), T + 5minutes)

specifies that if the gas has been turned on at a time T, but neither then, nor 5 minutes later the hob has been lit then the gas should be turned off.

tell(X, Y, request(R), T), have(Y, R,T) $\rightarrow$ tell(Y, X, agree-to-give(R),T1), T $<$ T1,
T1 $\leq$ T + 5

specifies that if Y receives a request for a resource R from X at time T, and Y has R at time T, then Y must answer within 5 time units with an agreement to give R.

Several proof procedures are available for abductive logic programs. For example CIFF [Endriss et al. 2004; Mancarella et al. 2009] is a proof procedure for ALPs that also incorporates constraint handling as in constraint logic programming. The constraint handling is used, in particular, for dealing with temporal constraints, for example in planning [Mancarella et al. 2004].

In Sadri [2007], preliminary suggestions have been made to use temporal abductive logic programming theories for specifying rules for assisting elderly people in their homes. Specifically mentioned are diary reminder and medication assistance systems. In the former, an ALP theory specifies when and what to remind the user about, for example to remind of a forthcoming appointment with the optician and that all spectacles have to be taken to the appointment. The reminders are issued according to the preferences of the user who can specify how frequently and how close to the event they wish to be reminded. In the medicine assistance system, it is assumed that pill bottles have sensors that can monitor the time when a dosage is taken out of the bottle. On this basis abductive logic programming is used to issue reminders of when the next dosage is to be taken, according to the time the last was taken and the required frequency. Also, in reaction to a medicine being taken, reminders are issued about any dietary restrictions that have to be followed, for example no tea or coffee to be taken for the next two hours.

4. AMI APPLICATION: HEALTH CARE

Existing eHealth applications have already improved the quality of health care. For example physicians can view radiological films and pathology slides in remote sites, and assist or perform surgeries via remote robots [Riva 2003]. Also there is a broad range of available sensor technologies to measure various respiratory, biochemical (e.g. glucose levels), and physiological (e.g. ECG) parameters. Gouaux et al. [2002] report the development of an intelligent wearable personal ECG monitor (PEM) as part of the European EPI-MEDICS project.[6] The motivation for PEM is early detection of cardiac events, improved decision making, and reduction of time before treatment is offered. PEM monitors the ECG, detects anomalies, and generates different alarm levels, depending on the perceived severity of the episode, and forwards them to appropriate care providers through wireless communication. There are three alarm levels, including: Major, indicating severe event, resulting in PEM directly sending a message to the emergency call center, and Low, resulting in advice being displayed to the user on the PEM screen to make an appointment with the attending physician.

AmI technologies can bring additional benefits by integrating and combining biological and physiological data with other data about the individual and providing further support within normal daily life [Haux 2006]. This is a new upcoming field of medical informatics, combining and developing such technologies together with information and communication technologies, to enhance environments for diagnosis, therapy, prevention, and early detection of diseases, as well as to support independent lifestyles of patients with chronic ailments. The emphasis of health care shifts from cure to prevention [Haux 2006].

Currently health care ICT (information and communication technology) is considered highly relevant [Haux 2006]. In the UK, prior to the current economic downturn, a budget of 9.2 billion Euros had been assigned for improvements to the national computerized health information systems (HIS) over a 10-year period. The aim of such HIS is to contribute to the quality and efficiency of patient care and medical, nursing, and administrative and management tasks and needs. This managed care focuses on organizational and management controls to provide cost effective appropriate care, and to ensure that health care remains affordable to populations with significantly lowering PSRs (potential support ratio, the number of persons aged 15–64 years per one person aged 65 or older).

Ubiquitous computing health-enabling technologies include wearable devices, such as microsensors embedded in textiles and personal computers, with wrist mounted displays worn on belts. These developing technologies are aimed at making it easier for individuals to monitor and maintain their own health while enjoying lives in normal social (noninstitutional) settings. The hope is that such enabling technologies will allow better quality of life into old age, beyond traditional health care, and also save money by reducing the need for institutional care [Haux 2006].

[6] http://epi-medics.insa-lyon.fr/statico/epimedica.htm.

AmI technologies may allow continuous monitoring in the home environment, which can lead to early recognition of deteriorating health and behavioral patterns. Some patterns, affecting health or indicating early signs of physical or cognitive deterioration, may not be easily detectable in a clinical environment, and may only be detectable through monitoring at home. For example someone may fail to administer their medication correctly, but only when the television is on or when they have a visitor, or someone may repeatedly go in and out of a room when the phone rings before answering it. Thus AmI promises to make a threefold contribution to eHealth, namely to integrate the fragmented systems into a connected system [Fifer and Thomas 2002], to provide personalized care, and to focus on preventative care [Riva 2003, Haux 2006].

Riva [2003] provides a motivating scenario as follows. Mario is an obese 40-year-old with type 2 diabetes who is being provided with a personalized self-management education program. The main features of the scenario are: recognition of Mario and his personal area network as soon as he enters the hospital, a micropayment system that automatically transfers the visit fee into the e-purse of the hospital, the fast provision to Mario via mobile devices of the diagnostic tests he needs to undergo and the location of all the specialists he needs to see on his schedule. All advice already given to Mario and his personal data, biomedical readings, personal habits, and so on, are available to the specialists, who in addition, can track the location of Mario and all their other patients on their office monitors. In the case of any problems or delays, the schedule of Mario's visit is modified with the aim of reducing his waiting time, and the new schedule is relayed to him via the same mobile device.

Haux [2006] argues that in addition to trends towards individualized and more preventative health care, with current and future developments in transinstitutional health information systems, better patient-centered health provision can be extended from regions, to nations, and around the globe. This is in line with the WHO (World Health Organization) eHealth resolution in their 2005 World Health Assembly.

The resolution urged each member state:

—to endeavor to reach communities, including vulnerable groups, with eHealth services appropriate to their needs;
—to develop the infrastructures for information and communication technologies for health as deemed appropriate to promote equitable, affordable, and universal access to their benefits, and to continue to work with information and telecommunication agencies and other partners in order to reduce costs and make eHealth successful.

There are, of course, risks attached to such developments, the primary one being the security of data. In these envisaged pervasive health enabling systems, the amount of data available on individuals may be significantly greater than in traditional health information systems. This data, together with its greater integration and persistence, are some of the main features that promise to help provide the individualization of care, but they are also features that increase the risk and damage of disclosure and misuse.

In the following we review several papers that have explored agent technologies, learning, and case-based reasoning techniques in applications of ambient intelligence in healthcare, and in particular in hospitals and residential care institutions.

4.1. Munoz et al. [2003], Favela et al. [2004], Rodriguez et al. [2005]

These papers report studies conducted in a hospital, which in collaboration with the hospital workers, led to the development of three scenarios for information management and interaction in that environment. The scenarios were then used for the specification of autonomous agents for the design of ambient intelligence environments for hospitals, and also for the design and implementation of an agent-based middleware, called SALSA.

The broad features of hospitals relevant to their study and reflecting their use of agents are the distributed nature of information, the need for collaboration, mobility of the personnel and devices, and the need to access accurate medical information in a timely fashion for decision-making.

One scenario involves a doctor (Garcia) who is examining a patient in a ward and receives a message on her PDA with a floor map indicating that the X-ray results of a patient in a nearby bed (patient P in bed 225) are available. She moves to a nearby hospital information system (HIS) display, which detects and identifies her and changes the display to one personalized to her, including her messages, personalized calendar, and information about the patients she is in charge of. The doctor selects the information about patient P and views the recent X-rays. The HIS, which is aware of the context, that is the patient, the newly arrived X-rays, and a possible need for a diagnosis, opens a window with access to various items of information, including the hospital medical guide relating to P's current diagnosis, and references to previous similar cases. Garcia selects to view a previous patient's records. The system infers that Garcia is interested in finding information about alternative treatments for P and recommends Web sources of digital libraries considered reliable by the hospital practitioners.

Another scenario involves a hospital map display alerting a cardiologist to the presence of a traumatologist on the next floor. Opportunistically the cardiologist sends a message to the traumatologist requesting a discussion session regarding a patient. The traumatologist accepts via PDA, approaches a HIS display unit by means of which the two specialists have a collaborative session about a patient's case while accessing information about him online.

The third scenario involves a patient, Ana, in a waiting room, who requires her PDA to recommend a cardiologist. The PDA makes a recommendation and in addition informs Ana that one of the recommended cardiologist's patients is in that waiting room, in case Ana wishes to approach her for a conversation.

In the SALSA architecture, agents are used as abstractions, to act on behalf of users, to represent services, or to provide wrapping of complex functionality to be hidden from the user. The system is required to be context-sensitive, where context is identified by user and device identification and location, role of users, time of events and interactions, and users' personal information stored in their PDAs. The agents have a standard life cycle of learning, through observations and messages (all through XML messages), thinking, and executing actions, including communications.

A hospital prototype system has been implemented. This includes a context-aware client, and a hospital information system (HIS). For the purpose of the former, doctors and nurses carry handheld computers that estimate their location, and depending on their roles, the client can inform them of other users, nearby, in the form of a displayed list. The HIS records patients' data, and the HIS agent monitors the changes in such data, and uses rules (much like ECA rules) to decide what information should be communicated to which users. For example when the HIS agent becomes aware that a doctor is in the vicinity of a patient and the patient's test results have been entered, the agent informs the doctor.

Rodriguez et al. [2004] describe in greater detail, the approach taken in location tracking in the hospital application and report related experiments. The approach is based on the use of radio frequency signal strength between mobile devices and access points of a WLAN (wireless local area network) infrastructure. A signal propagation model can be used to estimate the distance between the mobile device and the access point on the WLAN. But greater sophistication is needed to deal with complications such as walls, furniture, and other people in the way. To address this, Rodriguez et al. [2004] use neural networks, and a back-propagation learning algorithm. The neural network is trained by a variety of signal strengths at several locations in a building, for the four directions, north, south, east, and west. The location estimation component is integrated within the SALSA architecture by wrapping the trained neural network as one of the agents in the architecture. In the hospital application, given the decay of the signal strength from floor to floor one location estimation agent was trained for each floor.

4.2. Kofod-Peterson and Aamodt [2006]

This is another work addressing context awareness applied to a hospital environment. This paper, in follow-on work from an EU (Framework 5) project called AmbieSense, focuses on a medical domain with the aim of supporting health workers participating in diagnosis and treatment of patients. It uses case-based reasoning (CBR) to provide context classification.

CBR (see, for example Schank [1982]) is a problem solving technique that stores problems with their known solutions. Then, given a new problem, it finds a case from the recorded cases whose problems best match the new one. Then to solve the new problem the solution of the old matched case is adapted to the new problem.

In Kofod-Peterson and Aamodt [2006], a context is based on several parameters, including the location, the identity and role of people present, the time of day, and any artifacts that are present, such as patient list or patient chart. Each context has an associated goal. For example, if the context is "preward round" its associated goal is to "evaluate treatment of the patients."

The system architecture is as follows. There are three layers.

—The *perception layer*, which, collects data about the environment regarding the parameters that are used to classify the context;
—The *awareness layer*, where case-based reasoning is applied to match the data from the perception layer against the stored cases to find the closest match. The matching then also identifies the goal of the observed situation.
—The *sensitivity layer*, which, receives the goal recognized by the awareness layer together with what is known, through the perception layer, about the artifacts present. The sensitivity layer has, for each goal, a sequence of tasks that will achieve the goal. So given a goal and a possible list of artifacts it decomposes the goal into subtasks. The list of artifacts can aid the sensitivity layer to choose between two alternative decompositions. Each subtask is associated with an application agent that corresponds to an artifact in the recognized context that can provide some services.

The model was instantiated through observations at St. Olav Hospital in Trondheim in the summer of 2005, where a student followed various health workers and recorded the activities and events. Two wards were studied, cardiology and gastroenterology, to provide the cases. In this application, a case is a meeting between health workers, possibly also with patients present. Each such meeting has a purpose that provides the goal of the case. Around 200 cases were stored.

An example case is a preward meeting in which the duty physician and nurse participate, and it takes place in the morning. Its goal is to discuss treatment of the patients on the ward. This can be the best match for an observed situation of a meeting taking place in the morning in a doctor's office between a nurse, and another person who has a role of patient care. The observations can also include the data that the artifacts present are the patient list, patient chart, and PAS (patient administrator system). The goal of the meeting is then decomposed into subtasks, including: acquire name of patient from patient list, acquire changes in patient's condition since yesterday from nurse or patient chart, acquire any new test results from patient chart or PAS, and note any required changes in treatment (nurse or PAS).

4.3. Corchado et al. [2008]

This paper also reports work based on case-based reasoning and agents carried out on location at a care institute. It describes a system called GerAmi (Geriatric Ambient Intelligence) and its prototype implementation and testing at a care facility for Alzheimer patients. The care facility is the Alzheimer Santisima Trinidad Residence in Salamanca, Spain, which provides residential care for a maximum of 60 patients, all over 65 years old. The objectives of the GerAmi system are to monitor the patients and manage the work of the doctors and the nurses.

The system contains ID door readers, one above each door in the facility and the door of the elevators, ID bracelets for the patients and the nurses, with each bracelet containing an RFID chip, PDAs for the nurses, controllable alarms and locks, and wireless access points. The system architecture uses a multiagent structure with agents using BDI[7] (belief, desire, intention) concepts and case-based reasoning and planning. The implementation incorporates 30 patients, 10 nurses, and 1 manager, each of whom is associated with an agent. Further work is planned to incorporate 2 doctors with their associated agents. The manager and the patient agents run on a central computer and the nurses' agents run on mobile devices.

The patient agent records the location of the patient hourly and sends the record to a central database. The manager agent plays two roles. One is to control locks and alarms based on patients' locations, and the other is to allocate tasks to the nurses. For the latter it uses information about which nurses are available, taking into account holiday rosters, and which patients have to be served and what services (tasks, such as cleaning, feeding, exercise) are required. It also uses the profiles of the nurses, in terms of their expertise in providing services, and the constraint that they should not work more than 8 hours a day. For each task, it uses an estimate of how long it will take based on information about the maximum amount of time the best nurse for the task will take across the patients. The task allocations, in the form of a timed schedule of tasks, are then transmitted to the nurse agents via the nurses' PDAs. The nurse agents then make plans for the allocated tasks using case-based reasoning, via a library of plans. Replanning may be necessary, for example if the patient is not in the right location, or a resource is not available, or if an emergency occurs. Plan modifications are recorded and stored in the knowledge base of the cases, thus providing a learning capability.

Experimental results were collected over a period of three months. They pointed to three main conclusions. One is that GerAmi had the effect of reducing the time nurses spent on what the paper calls "indirect-action tasks,"—tasks not directly involved with the patients, such as monitoring, providing reports and visits. Another is that during the period of testing there was a reduction in the number of nurses working over 24 hour periods, even though the number of patients remained the same. Finally, as the system progressed, and thus the cases grew, the amount of replanning needed was reduced substantially—a figure of 30% reduction is reported in the paper.

[7]BDI [Rao and Georgeff 1991, 1995] is an agent model that deals with beliefs, desires, and intentions. The details are unnecessary for our purposes here, but in GerAmi the beliefs include knowledge about the whereabouts of the patients and the expertise of the nurses, desires are tasks the nurses have to perform in given circumstances, and intentions are the detailed plans of the nurses for performing these tasks. In GerAmi, as in the BDI model, planning is done using plan libraries.

5. AMI APPLICATIONS: SHOPS, SHOPPING, RECOMMENDER SYSTEMS, BUSINESS

This section includes papers on a variety of topics related to shops and business. We look at shops as responsive environments, with devices controlled by software agents that react to the presence of customers according to the customers' identities and profiles. Then we look at a proposal for ubiquitous commerce, which brings e-commerce and AmI together, with a framework for context-dependent interaction between shops and shoppers. We look at work on recommender systems in open connected environments, with emphasis on adapting known user profiles in one domain for use in another domain. The section continues with a proposal for AmI-based product aftercare and maintenance, and concludes with some example demonstrations of AmI-based products, including fabrics with embedded electric units.

5.1. da Silva and Vasconcelos [2005, 2007], Shops as Responsive Environments

These papers describe how software agents can be used to model and implement responsive environments. A responsive environment is described as one which senses events occurring in it and reacts to these events. This reactive behavior is provided by software agents, and the responsive environment is a combination of a physical world and agents.

The architecture described is generic, but the paper focuses on the application to an AmI-enhanced bookshop. The bookshop is equipped with sensors distributed among the bookshelves and throughout the store. These can detect and identify customers from their portable devices, for example, Bluetooth-enabled mobile phones, PDAs, or store loyalty cards. The store has some information about the profiles of its customers, such as what they have bought before and their preferences, and it can associate the customer identity with his/her profile.

The goals are to use all this information to provide customers with a pleasant experience while they are in the store and to encourage them to buy, in order to maximize sales. To these ends the store has various devices, for example there could be music-playing devices and LCD displays scattered throughout the store. The music devices should take into account the preferences of the customer closest to them when choosing what music to play, and the LCD displays should show books or other items that are likely to be of interest to customers closest to them. Similarly, the lights in the shop display windows should highlight the items that may be of interest to the nearby customers. Moreover, adjustments should be made when multiple customers are nearby, with a smooth transition of displays to serve interests of all customers within the range.

The architecture chosen for this application is that of autonomous communicating agents. Each customer entering the system in and within some outside range of the shop and each device in the shop has a corresponding agent associated with it. Figure 2 is a representation of the diagram given in the paper. The rectangle is the shop environment with customers (the smiling faces) and devices (d1, d2). The cloud is the software where each element of the shop environment is mapped onto an agent. Within the shop, each device has a range, represented by the oval enclosures. The arrows connecting the agents represent communication among the agents. This communication is via blackboard style tuple spaces, providing ad hoc communication, implemented via Linda tuple spaces in Prolog. The system also includes administrative agents, not shown in the cloud.

Administrative agents roam the system and continuously check the tuple space for tuples that concern them. When a component, for example, customer, enters the system and is identified, its ID and location are entered on the tuple space, and an administrative agent will start up an agent for the component, if one does not exist already. Changes of location are likewise entered on the tuple space. The agent that

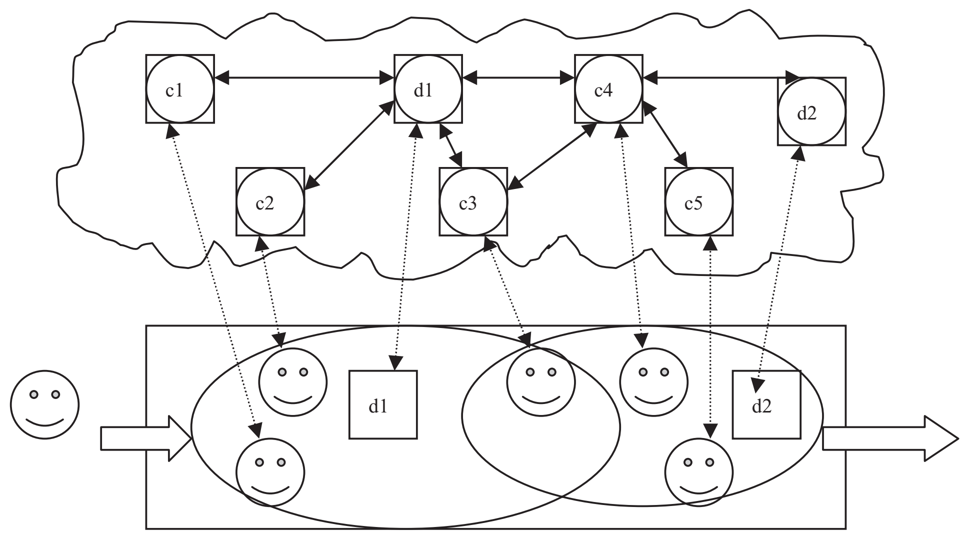

Fig. 2. An architecture for agents in a responsive environment.

is started up functions according to an observe-process-act type cycle implemented in Prolog, where observing and acting mean taking tuples from, and putting tuples on, the tuple space, respectively. The agents' notion of environment is the tuple space. The agents have goals that are implicit within their observe-process-act programs.

In the bookshop scenario, say, a customer's presence is detected within range of a display device. The customer's agent writes the customer's book type preference on the tuple space (e.g <price, age> = <20, 20>, meaning customer is 20 years old, and price maximum is 20). The agent associated with the device takes this and displays items corresponding to this preference (assuming that the bookshop has classified its items according to these two numerical parameters).

This procedure is modified to cater for the presence of multiple customers within the range of a device. The idea is to smoothly move from the displays of items of interest to one customer to items of interest to another customer. To achieve this, each customer agent has a pair of parameters $<\Delta p, \Delta a>$, with fixed values, Δp, $\Delta a \in [0,1]$; Δp is associated with price and Δa is associated with age. These $<\Delta p, \Delta a>$ parameters appear to be assigned to customers arbitrarily. If there are multiple customers within range of a display device, their agents take turns to update the <price, age> tuple on the tuple space, thus influencing what is displayed, according to the following instructions:

$$\text{Replace the current } \textit{price} \text{ by: } \quad price + \Delta p^{*}(price_c\text{-}price)$$
$$\text{Replace the current } \textit{age} \text{ by: } \quad age + \Delta a^{*}(age_c\text{-}age),$$

where $<price_c, age_c>$ is the preference of the customer represented by the agent updating the tuple space.

5.2. Masthoff et al. [2007]

This paper reports continuation of the work reported in da Silva and Vasconcelos [2007], and summarized in Section 5.1 but with most of its focus on catering for multiple customers within the range of devices. The scenario is the same here, namely a bookshop equipped with sensors to detect the presence of customers, and display and audio devices distributed throughout the store. The architecture is also the same, namely a multiagent system, with agents representing customers and devices, and communicating via a blackboard mechanism. Here, a proof-of-concept implementation is reported, which uses JADE for the agents and JavaSpaces for their communication. An off-the-shelf Bluetooth USB adaptor is used to detect Bluetooth-enabled mobile phones. The owners of the mobile phones register their likes and dislikes in terms of music genres and artists and these profiles are stored in a database accessible by the software agents.

The paper considers several existing strategies for aggregating ratings of individuals into a group rating. The objective is to explore ways of deciding what excerpt of music to play next, for example, when several people are near the device and their individual preferences are known. The strategies are inspired by social choice theory, and include the average strategy, which takes the average of the individual ratings, and the least-misery strategy, which takes the minimum individual rating.

One assumption is that the membership of groups of people changes continuously, and another is that the members are unrelated (not families or friends), and in particular they do not affect each other's preferences (no contagion effect). Experiments with human subjects are reported briefly, and one result used in the implemented satisfaction function is that how much an individual has liked what has been presented so far affects how much they will like what comes next. The satisfaction function, modified from Masthoff and Gatt [2006], evaluates the satisfaction of seeing an item i, based on the impact of i and the satisfaction of seeing a sequence of items prior to i:

$$Sat(items + <i>) = (\delta^{*} Sat(items) + Impact(i, {}^{*}Sat(items)))/(1 + \delta)$$

with Impact(i,s), the impact on satisfaction of new item i, given existing satisfaction s, is defined as

$$Impact(i, s) = Impact(i) + (s - Impact(i))^{*}\varepsilon.$$

Here δ represents the decay factor of previous experience; when it is 0, past items have no effect, and when it is 1, there is no decay and past items have maximum effect. Intuitively, δ represents the degree to which previous experiences can affect the new experience. Experimental results of the system are not reported.

5.3. Keegan et al. [2008], Easishop

This paper also focuses on applications of AmI in shopping, with the aim of providing a synergy between AmI and e-commerce to bring about what they call U-commerce, namely ubiquitous commerce. Their work is within the context of the Single European Electronic Market (SEEM) [SEEM 2003] vision, which has the objective of developing an electronic framework in which different participants in the economy can collaborate. The authors describe the architecture and implementation of a system called Easishop, which aims at providing context-dependent exchanges of information and negotiation between shops and shoppers. The architecture is based on a multiagent structure with the agents equipped with BDI (belief, desire, intention) concepts. Each shopper has his own agent, as does each shop.

A scenario described in the paper is as follows. A shopper requires boots of a particular make and size, and may have price constraints. This is input in the shopper's agent. As the shopper approaches a shop in a shopping centre, the shop's hotspot scanner spots the shopper's agent, and the latter detects that it is within the shop's hotspot. The shopper's agent then migrates to the shop's Easishop node, checks if the shop has the required item and negotiates a price. Depending on the outcome the shopper's agent may decide to initiate an auction among the shops in that area. To do this the agent migrates to the Easishop e-marketplace and advertises the required item. The e-marketplace agent, then invites the shops within its range to a reverse auction, collects the data, including the prices they offer and the shops' distances to the last known location of the shopper, and passes it on to the shopper's agent. This agent returns to the shopper's device (e.g. smartphone), adjusts the information about distances to shops, as the shopper may have been moving while the agents have been active, and presents the information to the shopper. The rest is then up to the shopper, to chose if he wishes to go ahead with a purchase and if so, from which shop.

A prototype implementation is reported covering much of this scenario. Communication is via Bluetooth short-range wireless. The shopper's devices are mobile phones (SonyEricsson P910i) and PDAs (HP IPAQ 3870). All software is implemented using Java, and the interface uses a GUI toolkit called Thinlets (http://thinlet.sourceforge.net/). The agents are implemented using a framework called Agent Factory [Collier et al. 2003]. The shopper inputs his profile (e.g. sex, age, shoe size) into his device. He can add items to his shopping list by navigating a product list presented as a tree structure, first identifying the product class and then becoming more specific, in effect navigating a tree rendering of an XML schema of the UNSPSC (United Nations Standard Products and Services Code)[8] set of production description codes.

The agents behave according to the commitment rules that their designer endows them with. Commitment rules in Agent Factory specify what commitments the agent should adopt given its current beliefs. For example, the following commitment rule fires when the communication agent associated with the shopper device detects a shopping hotspot. As a result of firing, the agent drops a commitment to scan for hotspots, and informs the shopper's agent of the hotspot. $SEQ(A_1, A_2, \ldots, A_n)$ represents the sequence of actions $A_1; A_2; \ldots; A_n$.

```
BELIEF(connectedTOHotspot(?hotspotname)) =>
COMMIT(Self, Now, BELIEF(true),
    SEQ(retractBelief(ALWAYS(BELIEF(scanning))),
Inform(agentID(DeviceAgent), connectedToHotspot(?hotspotname)));
```

The following two example commitment rules (slightly modified from the paper for easier reading) result in the shopper agent first adopting the belief that it is in a hotspot, when it is informed of this, and then migrating to the store associated with the hotspot if it believes that the shopper wants a product. The term *fipaMessage* is a reference to a FIPA-compliant message.[9]

```
BELIEF(fipaMessage(inform, sender(?name),
    connectedToHotspot(?hotspotname))) =>
COMMIT(Self, Now, BELIEF(true),
adoptBelief(ALWAYS(BELIEF(connectedToHotspot(?hotspotname)))));

BELIEF(shopperAgent(?shopperagent)) &
BELIEF(wantProduct(?description, ?ean, ?urgency, ?range, ?location)) &
BELIEF(connectedToHotspot(?hotspotname)) =>
COMMIT(Self, Now, BELIEF(true),migrateAgent(?shopperagent));
```

User evaluation is part of ongoing and future work.

[8]http://thinlet.sourceforge.net/.
[9]FIPA stands for The Foundation for Intelligent Physical Agents, http://www.fipa.org/, and it is an IEEE Computer Society standards organisation for agents and multi-agent systems.

5.4. Gonzales et al. [2004a, 2004b, 2005, 2006], Recommender Systems

Several papers by Gonzales et al. [2004a, 2004b, 2005, 2006] discuss ambient recommender systems, where the focus is to provide personalized information without increasing the need for direct feedback and input from users.

Gonzales et al. [2004b] focus on capturing user data for recommender systems. The primary aim of their proposed system, called Smart User Model (SUM), is to lower the burden on users for initialization procedures involved in providing data, while still aiming to provide users with personalized information. SUM aims to capture user information in a generic way in order to make it feasible to transfer user data from one domain in which the user has been profiled to another in which he has not. The assumption is that the next generation of recommender systems will have a portable user model to interact with multiple open, distributed, heterogeneous services, and will use ontologies to transfer user preferences across the services.

The following example scenario is given. The user, say John, has interacted with a restaurant and a cinema recommender system, and these have captured John's profile, including objective information, such as name, age, sex, country; and subjective information, such as the information that he likes attractive places and imaginative cuisine, but efficient service is not very important to him. A marketing recommender system has not interacted with this user, but the aim is to adapt the profiles already captured by the restaurant and cinema services to allow the marketing service to also make appropriate personalized recommendations.

The user models are divided into 3 categories, objective, for example age, sex, country; *subjective*, for example preferences, and emotional, which is not addressed in these papers. The objective attributes can be provided by the user and obtained via a database. The subjective ones have to be acquired through user interactions with the environment. The subjective values are represented by values in the interval [0,1], each value representing notions such as a little bit, very, very much, and not. So, for example John's subjective information regarding the restaurant may be represented as {<attractive-place, 0.7>, <imaginative-cuisine, 0.8>, <efficient-service, 0.1>}, specifying that imaginative cuisine is very important to him, much more so than efficient service.

Mappings are suggested, that in combination may allow the transfer of the user profile in one domain to another. In effect, in the case of objective attributes, essentially, to transfer to another (new) service the authors combine all objective data that is known and intersect the categories with those required in the new service, assuming uniform ontologies. In the case of subjective attributes, the mapping is done in two stages. First the known subjective profiles related to known domains are mapped onto a more generic SUM profile. For example, John's SUM profile may be {<attractive, 0.7>, <imaginative, 0.8>, <efficient, 0.1>}. For the second stage it is assumed that a priori there are weighted pairings of some SUM attributes with the attributes of the new domain (service). Continuing with John's example, if the marketing domain has attributes including novel, useful, and creative, then a priori pairings may be assumed between, say, attractive and useful, and between imaginative and creative, with the weighting providing some indication of confidence in such pairings. Thus a subjective profile is constructed for the new domain with attributes inheriting values from their SUM counterparts, but with added weights representing confidence.

5.5. Kopacsi et al. [2007], Weber [2003], Product Service and Aftercare, Business

Kopacsi et al. [2007] discuss the possible role of AmI in manufacturing. The authors argue that economy and environmental requirements are imposing new trends on manufacturing, whereby not just products but long term product aftercare and services will be in demand, and subject to competition. Thus, Kopacsi et al. argue, there will be increasing demands for product life-cycle management (PLM), and in particular, conditional and predictive maintenance, and it is in these that the contribution of AmI may be important. More conventional forms of maintenance are often either based on breakdowns, where repairs are done when something breaks down, or are based on periodic and systematic checks and repairs. But if sensors can provide continuous or regular information about products and this information can be integrated with background information, for example the history of the product or similar products, then intervention can take place before a problem arises, and only when needed.

The paper describes work on a European project called FOKsai, which is concerned with providing knowledge management tools and AmI support for SMEs (small to medium enterprises) to support extended product management and aftercare. The project has the participation of four European SMEs and the goal is to provide knowledge bases and diagnostic tools to make the aftercare provided by the companies more efficient and more effective.

For each SME, the system includes a knowledge base, which is based on the relational model using Oracle and MySQL. This is a repository of product-related knowledge, and the AmI-related data provided by sensors which is translated to tuples and input into the database. We will see later (in Section 7.1) that this is also the approach taken by Feng et al. [2004]. There is also a diagnostic engine that uses case-based reasoning[10] implemented in C++. This is used to help the diagnosis of problems and provide solutions, and is updated as the system evolves. The system allows remote access and provides suggestions on how to proceed in reaction to ambient information.

One SME produces cutting machinery, and provides 24 hour customer support, often involving site visits by maintenance engineers to solve breakdowns. The proposed system aims at cutting these overheads by continuously monitoring the operation environment, for example the temperature, and automatically warning if corrective action needs to be taken. For example if the temperature falls too low the system informs the operator that the cutting should be stopped and temperature raised, otherwise the quality of the cutting will be poor. Another SME provides networked IT security and safety systems for a multitenancy building. It needs to maintain a network of heat, motion and other sensors and is hoping to cater to the needs and preferences of the individual tenants in the building.

A different project, Werner Weber [2003] summarizes several industrial projects in AmI at the Infineon Lab for Emerging Technologies in Germany. Some involve the integration of electronics into fabrics, such as clothing, textiles for floor covering, or bed sheets. One demonstrator is a jacket with an integrated MP3 player with a key pad, ear phones, microphone and battery. The integrated devices are left inside the clothing when cleaning. Another demonstrator is a device that converts body heat into energy. The energy produced is small, enough to power a wrist watch or a heart rate sensor. A third demonstrator is fabric that has a distributed network of electronic units woven into it. Such fabric could be used in carpets and other floor coverings for surveillance or guidance in buildings, or in intelligent sheets that monitor the vital signs of hospital patients.

6. AMI APPLICATIONS: MUSEUMS, TOURISM, GROUPS AND INSTITUTIONS, OTHERS

In recent years the tourism industry has realized the advantages of providing personalized tailor made information to users [Rumetshofer et al. 2003]. This section starts by describing work in that direction, first in the context of museums and heritage sites and then in the context of visiting a city. We then look at contributions of AmI in

[10] See section 4.2, Kofod-Peterson and Aamodt [2006], for a brief description of case-based reasoning.

group decision making support and in institutions, and end with an AmI application in monitoring driving.

6.1. Busetta et al. [2003, 2004], Penserini et al. [2005], Museums

The work described in Busetta et al. [2003, 2004] concerns responsive/active environments, and in particular interactive museums. The work is part of the project Peach (Personal Experience with Active Cultural Heritage[11] which is funded by a national body in Trento, Italy, and involves collaboration between academics in Italy, Germany, USA, and Canada, and heritage organizations in Italy (Castello del Buonconsiglio, Pompei Excavations) and Germany (Volklinger Huette Museum).

A typical scenario considered is one in which a visitor to the museum requests, possibly via his PDA, a presentation about one of the exhibits. Several facilities, modelled as agents, are able to produce the presentation with different capabilities, for example pictures, audio, video, or a combination of audio and video. The visitor is near some facilities for the presentation, for example screens or speakers, and his PDA may also have the capability of displaying the presentation in a limited way. The visitor may be part of a group, for example a family, or in the vicinity of other people, and the museum may have information about the visitor's interests, for example from previous visits.

The architecture chosen is a multiagent system capable of forming implicit organizations. The organization is implicit in that it is not preprogrammed and there is no explicit formation phase. This is facilitated by role-based communications and overhearing. The experimental communication infrastructure, called LoudVoice, involves streams of messages (called multicast channels) that can be heard by many agents. The FIPA-compliant messages have headers such as REQUEST, QUERY, INFORM, DONE, and the senders and receivers can be role identifiers, rather than individual agent identifiers.

An implicit organization is a set of agents that play the same role, for example Presentation Planner. These agents may change dynamically and may have different capabilities and there may also be some redundancy among them. Agents may have several coordination policies, for example Plain Competition, whereby all attempt to provide the service requested and the first to do so wins. Another is Simple Collaboration where there is some degree of synthesizing of the answers obtained independently by individual agents. When a request for a presentation is sent, the agents currently available with role Presentation Composer implicitly form an organization. In this organization, a coordination policy is chosen among those that are common to all the agents involved. Then each available Presentation Composer agent may request various items of information from agents playing other roles. The information includes user information (for example location of user) from agents playing the role of User Modeller, presentation data (e.g. data about the exhibit) from agents playing the role of Information Modeller, and display information (e.g. free available display medium compatible with the capabilities of the composer agent) from the agents playing the role of User Assistants. When all the information has been received, the Presentation Composer organization can decide its answer based on its chosen coordination policy.

Further related work [Penserini et al. 2005] uses the Tropos methodology [Bresciani et al. 2004] to provide a model for the implicit organization architecture, in general, and example instances of it for the active museum application, in particular. Tropos is a methodology that adopts agent concepts for requirement analysis and software engineering. In particular it uses concepts such as actors, agents, roles, goals, tasks, resources, and different types of dependencies among actors. The dependencies between one actor and another, can for example, indicate that one delegates fulfillment of a goal

[11]http://peach.itc.it/.

to the other, or that one is required by the other to perform a task, or that one is required to provide a resource to the other.

6.2. Costantini et al. [2008], Heritage Sites

This paper describes the use of multiagent technologies in combination with sensor and satellite tracking and PDA technologies for two heritage applications. One application is in providing personalized information for visitors to Villa Adrianna, which is an extensive ancient Roman villa. The other application is in tracking heritage artifacts during transportation, to prevent theft and to maintain an ambient atmosphere in terms of temperature and humidity. The multiagent system used, called DALICA, relies primarily on logic programming for its specification and implementation.

For the Villa Adrianna application, the system uses a model of the site in term of its points of interest (POIs). Each POI is represented as a fact giving its location and keywords characterizing its main features, such as garden, water, statue, mosaic. Each such feature is expressed with an associated weight that signifies the prominence of that feature at that POI. For example the fact

poi('VA_IlPretorio', 41.939503, 12.775775, 25, [(columns, 0.30), (opus, 0.30), (fresco, 0.30), arch, 0.10)],8)

denotes the location of the Villa Adrianna Pretorio as a circle with center at (41.939503, 12.775775), according to the satellite coordinates, and radius 25 meters, with the given list of primary features, columns, opus, fresco, arch, with the arch being the least prominent feature. It also denotes that it will take about 8 minutes visiting it. The visitor can provide DALICA with his initial list of interests according to the features used in the definition of the POIs. DALICA can also infer the user interests from the POIs he visits, the length of time he spends at each and the POIs' lists of features. The interests can be confirmed by the user. In addition DALICA monitors the visitors' passage through the heritage site for other purposes, for example to suggest POIs of possible interest in their vicinity, and to issue warnings if the visitors trespass to restricted areas. The communication with the user takes place via a PDA, and the location tracking takes place via a satellite.

For the transportation application, each package is equipped with a mobile device that periodically checks temperature and humidity and the position of the package. It also checks the actual route via which the transportation is taking place, given by satellite information, against the agreed planned route. Any anomalies will result in the sending of warning messages.

Both applications have been implemented, and systematic evaluation is part of ongoing and future work.

6.3. Petersen and Kofod-Petersen [2006], Tourists

This paper describes a scenario about guiding tourists through the city of Trondheim, inspired by one of the ISTAG AmI scenarios [Ducatel et al., 2001]. The scenario is intended to incorporate the use of AmI for the tourist and the use of virtual enterprises for the businesses in Trondheim that wish to enhance the tourist experience and at the same time enhance the reputation of the city as an attractive place to visit and increase their own trade. The project, Wireless Trondheim, is a collaborative effort between the local council and the university in Trondheim.

In this scenario the tourist arrives in Trondheim on a ship and on leaving the ship signs for the Trondheim experience (TE), which provides her with a PDA. The goal of the tourist is to see and enjoy the city in one day, and she can express this goal by selecting some key points to indicate her interests, for example medieval Trondheim, local food, and shopping. The tourist's goal can be planned for by the TE to produce a sequence of activities. In this, information is used from virtual enterprises (VEs) formed by city businesses and service providers. For example, there may be one such VE consisting of a museum, a concert organizer, a restaurant, and a taxi company, collectively offering a single package, of a guided tour of the museum, a visit to the shopping centre, lunch at the restaurant, a concert at the cathedral and transport back to the ship. The VEs may have access to the tourist's profile collected by the TE and offer their package only if it is compatible with the profile.

The TE may choose a number of such VE packages to present to the tourist, according to what it knows about the tourist's desires and goals. Once the tourist chooses one package, the TE guides the tourist through the activities. It can also use the location trackers to provide information about the architecture or history of the buildings, shops, and bargains that may be on offer as the tourist passes them in the city. TE may also have access to other TEs guiding other tourists from the ship and may give information to the tourist about who else from her group may be having lunch at the restaurant and attending the concert.

6.4. Groups and Institutions

Group decision making, and in particular, mixed-initiative group decision making, where humans and computer systems, such as artificial agents, collaborate, is also an area that is being explored within AmI frameworks. For example, an area increasingly requiring group decision support systems (GDSS) is health care where patients' treatment may involve several specialists, distributed in various departments [Karacapilidis and Papadias 2001]. Jonathan Grudin [2002] provides the following classification of digital technology in Group Decision Support Systems: preubiquitous (70s), supporting face-to-face meetings, protoubiquitous (90s), supporting meetings distributed in space, and ubiquitous (current), supporting meetings distributed in time and space. Emerging work also proposes to take into account the emotional state and past histories of the current participants.

Prakken and Gordon, [1999] provide an early paper on the specification and implementation of an automated mediation system for group decision making. The work described relates to the Zeno computer system, developed at the GMD Bonn [Gordon and Karacapilidis 1997] and applied to urban planning procedures, in the context of an EU-funded project called GeoMed. The objective is to monitor electronic meetings, and remind the participants what moves should be made next, and alert when rules are broken. But it is also required to be flexible enough to allow meetings to proceed even when rules are broken.

The project has chosen Robert's Rules of Order (RRO), which are well known rules, which capture standard procedures in the US for group decision making, based on US parliamentary procedures. In brief, the rules govern how and in what order, the floor is obtained, motions are obtained and seconded, debate is opened, voting is called for and the state of the meeting is updated.

The paper provides a formalization of RRO in first-order logic, augmented with some quasideontic features, using states and hierarchies of speech acts. For example,

x Is stated by chair at s

represents that the chair has uttered statement x in state s.

m Is a motion $\wedge$ *m Is seconded by p at s*

represents that motion m is seconded by p in state s.

Type: x Is a motion
Superclass: x Is an act

represents motion as a subclass of act.

$\forall x \forall s$ *(x Is stated by chair at s* $\wedge$ *x is debatable* $\rightarrow$ *x is open to debate at s')*

represents that a debatable motion becomes open to debate after it is stated by the chair. Here *s'* represents the state immediately following *s*.

$\forall x \forall y \forall s$ *(y Correctly makes x at s* $\leftrightarrow$ *x Is an act* $\wedge$ *y Makes x at s* $\wedge$ *x Is in order at s* $\wedge$ *x Is proper at s)*,

states a deontic feature that a procedural act is correctly made in a state if and only if it is in order and proper in that state.

A Strips-like approach is used to update the representation of the state of the meeting, as the meeting proceeds. The formalization allows recognition of actual versus required behavior in the meeting. It recognizes and reports violations of the required bahavior, but does not enforce the rules (e.g. by blocking voting, when the motion has not been seconded). If the chair overrides a rule and breaks it by an action the system records the action, updates the state and carries on.

More recent work [Marreiros et al. 2007a, 2007b] proposes an agent-based architecture to support ubiquitous group decision making. The objective of this work is to provide a system that can simulate group decision and can answer what-if questions. In their proposal each human participant can be associated with several agents, and it is assumed that the group is persistent, as the group's previous decision making processes and the members' past behaviors are to be used in the process. The intention is for the system to take into account past behaviors, current exchanges of dialogues, and information about emotions of the participants, to suggest arguments to be exchanged with other specific participants. An implementation is reported using Open Agent Architecture, Java and Prolog, but the details of the reasoning and knowledge representation are not given. The dialogues between agents involve requesting the other agent to perform an action, accepting or refusing a request, and requesting with an argument. An argument could be, for example, a threat, a reward or an appeal to past practice.

Boella et al. [2008] propose a theoretical framework for ambient intelligent systems based on the notion of institutions. As an application focus they consider a scenario of an enhanced classroom where the teacher and the students are provided with pocket PCs which allow for communications between students and between students and the teacher. The teacher has the added facility of being able to monitor the flow of communication between students and to stop it selectively if he judges that the students are being distracted from their work. Lessons conveyed via the PCs can be targeted, for example with more advanced material for the more advanced students, and discussion and other working groups can emerge naturally and virtually. The advantages of such a system could be that it allows distant learning, for example if a student is away or ill, and it allows the students to work at their own level, without the level becoming public. The teacher can ask questions on the system and collect answers. He can also collect feedback instantly to keep the lesson more focused at the correct level.

The proposed formalism starts with the agent's view, where functions associate goals, beliefs, and skills to agents, and another function associates skills for the achievement of goals. For example, with P representing the teacher and M, one of the students, *goals(P)* = {*obtain feedback on the topic of postmodernism*}, *skills(M)* = {*communicate with school friends, answer the questions*} represent P's goal(s) and M's skill(s), respectively. Also *rules*({*require feedback*}) = {*obtain feedback on the topic of postmodernism*} represents that the skill, *require feedback*, achieves the goal, *obtain feedback, on the topic of postmodernism*. Social concepts are then brought in via the notion of power and social dependence networks. These too are specified via functions; power is specified by associating with each agent, the goals it can achieve, and social dependency, by associating a set of agents with another set on which they are dependent for achieving a goal. For example *dep({M}, {P}) = {communicate with school friends}*, denotes that student M is dependent on teacher P for communicating with school friends (as the teacher has the power to allow or disallow this).

The formalizm is then augmented at the institutional level, in a similar way as the agent level, associating public goals, beliefs, and skills; this time not to individual agents, but to roles, for example role of teacher and role of student representative. No further theoretical results or analysis is given.

6.5. Bosse et al. [2008], Driving

This paper describes an agent model for monitoring driving behaviour. The objective is to take periodic sensor readings of the steering wheel operation of the driver and his gaze, and if over time it is detected that the driving is impaired, either to lock the ignition, if it is already turned off, or to slow and finally stop the car, if ignition is on. The idea is inspired by a commercial system currently under development by Toyota in their Lexus line.

The system proposed by Bosse et al. has four types of agents, sensoring agents, monitoring agents, a driver assessment agent and a cruise control agent. They can interact through one-to-one communication and can make observations and perform actions. All the agents are instances of the Generic Agent Model (GAM)[12] [Brazier et al. 2000]. The sensoring agents collect gaze information from the driver and steering wheel information from the car, and pass these on to the monitoring agents. The monitoring agents review this information over a period of time and verify if steering or gaze has become “abnormal.” If so, the information is passed on to the driver assessment agent, which can diagnose that the driver is in an impaired state, and communicate this negative result to the cruise control agent, which then either slows the car down or locks the ignition.

All the specifications are based on predicate logic, and are in some places, based on Temporal Trace Language (TTL) [Bosse et al. 2006]. TTL is a predicate logic-based language, including an executable sublanguage, incorporating states and time points. To give a flavor of the specification, we give some samples in the following. For example, the rules governing the behavior of the cruise control agent are as follows:

internal(cruise_control_agent) | belief(driver_assessment(negative)) ∧
internal(cruise_control_agent) | belief(car_is_not_driving) →
output(cruise_control_agent) | performing_in(block_ignition, car_and_environment)

internal(cruise_control_agent) | belief(driver_assessment(negative)) ∧
internal(cruise_control_agent) | belief(car_is_driving) →
output(cruise_control_agent) | performing_in(slow_down_car, car_and_environment)

The first specifies that if the cruise control agent believes the driver assessment is negative and the car is not being driven (is not on the move) then the car ignition is blocked, and the second specifies the car slow down action. All the agents share generic rules regarding beliefs generated by observation and communication, for example,

∀X,Y:AGENT, I:INFO_EL
input(X) | communicated_from_to(I,Y,X) ∧ internal(X) | belief(is_reliable_for(Y,I)) →
internal(X) | belief(I)

states that X believes information communicated to it by a reliable agent Y.

[12]The Generic Agent Model is a general purpose, domain-independent, model that conforms to the weak notion of agency, i.e. as well as autonomy and the capability of communication, it incorporates reactive and proactive behavior.

The paper presents a simulation run of the system. The simulation scenario involves the driver starting the car and driving. Then between time points 10 and 20 the sensoring agents record the steering behavior as:

output(driver) | performing_in(steer_position(centre), car_and_environment)
output(driver) | performing_in(steer_position(left), car_and_environment)
output(driver) | performing_in(steer_position(right), car_and_environment),

and the focus as:

output(driver) | observation_result_from(gaze_focus(far_away), driver).

The monitoring agents interpret these as indicating that the steering has becomes unfocused, and the driver's gaze fixed.

The following is an example of a "monitor property" rule that the monitoring agents use:

$\forall$t [t1 $\leq$ t $\wedge$ t $\leq$ t2 $\wedge$ belief(at(steer_position(centre),t)) $\rightarrow$
$\exists$t' t $\leq$ t' $\leq$ t + D $\wedge$ not belief(at(steer_position(centre),t')),

which states the property that between t1 and t2 whenever the steering is in a central position, then at a slightly later time (after D units, at most) it is not in a central position, that is the driver keeps on moving the steering. The verification of this rule by the monitoring agents leads ultimately to the passing of the information of driver_assessment(negative) to the cruise control agent which believes it, and consequently generates the action of slowing down the car.

7. AMI: OTHER DATA MANAGEMENT AND AI TECHNIQUES

In the previous reviews, we have seen several uses of AI, databases and agent technologies, including ECA-rules, temporal, abductive, case-based reasoning, learning and fuzzy logics. In this section we look at other AI and database technologies that have been explored and advocated for use in AmI applications. Among these we look at papers that identify the need for, and make proposals towards, context-aware data access, distributed planning, self-organization, and embodied systems.

7.1. Context-Aware Data Management

Feng et al. [2004] and Li et al. [2008] argue that in the past decades, work on databases has concentrated on content-based access, but AmI applications require context-aware access and data management strategies and solutions. Content-based data management focused on efficiency, whereas context-aware data management focuses on usefulness. A context-aware data management system should be able to deal with queries such as: "Get the report I prepared last night before dinner for this afternoon's meeting," or "Find restaurants nearby which I have not visited for half a year." Such queries will require different answers for different users and different answers at different times for the same user.

The authors consider context to be the situation in which the user tries to access the database. They identify two forms of context: user-centric, and environment-centric. Each can be divided into various subcategories.

A user-centric context may be:

—Background (e.g. interest, habit), from a user profile;
—dynamic behavior (e.g. task, activity), from a user agenda;
—physiological state (e.g. body temperature, heart rate), from body sensors; or
—emotional state (e.g. happiness, anger, fear), from multimodal sensors and analysis of user features.

An environment-centric context may be:

—Physical environment (e.g. time, location, temperature), from sensors;
—social environment (e.g. traffic jam, surrounding people), from service providers; or
—computational environment (e.g. surrounding devices), inferred from user location and activity.

The issues the authors identify include:

—How to acquire, categorize and model contextual information;
—how to exploit contexts to answer a user's data request;
—context-aware query language for users;
—how to effectively communicate answers to the users on small hand-held devices;
—what context-aware strategies are needed, both for finding useful answers to queries and for presenting the answers to the users.

They propose several strategies: (assuming a relational data model). Some are for modifying queries, for example:

(1) Context that adds query conditions based on the current situation: useful for queries such as "look for the earliest flight that I can catch," or "look for the fastest route to the airport given the current traffic conditions";
(2) context that adds query conditions based on recalling past events: useful for queries such as "what was the information I got in the hotel";
(3) context that adds conditions based on other factors:
   —User's real intention: for example when the user asks for information about nearby restaurants, the real requirement is information about nearby restaurants that are open now, as the user has to take clients to lunch according to his/her agenda.
   —Personalizing the query: for example the user asks for directions to drive to a destination. If it is night-time the preferred answers are those that avoid roads through dark woods.

Other strategies are concerned with how query results are returned to users, for example:

(1) For small devices sort query results so the most useful is given first. Another paper [van Bunningen et al. 2007] describes a technique for ranking query results according to the user's preferences, based on a probabilistic function of the user's past choices.
(2) Query result given in a way convenient to the present situation, for example, if the user is driving then give the query result as speech, if the user is talking to someone then postpone delivery and alert via a vibration, and if there is a group of people waiting for the result then present it on large screen.

The authors propose a two-layer context-aware data management architecture, consisting of a public data manager, namely any conventional database manager, and a private data manager incorporating the context-awareness. A user query is first processed by the private data manager, which uses profile, agenda, and log information to modify and augment the query, and then sends it to the public data manager for conventional query processing. The answer from the public manager is passed through the private manager and transformed for appropriate form and level of delivery and returned to the user.

Van Bunningen et al. [2006] provide further developments of this proposal to provide context-based preferences for database querying in AmI environments. In this

Table I. Person

Person
ID
Eric
Peter

Table II. hasActivityType

hasActivityType	
Source	Destination
Eric	Reading
Peter	Sleeping

work they distinguish between static preferences—preferences that are not context-dependent, and context-based preferences. Static preferences are catered for, simply by specifying which database tuples are preferred, for example "prefer the (tuple related to the) cheaper of two books with the same ISBN." Context-based preferences, on the other hand, are dynamic and situated, for example "Peter prefers TV programs of human interest genre when doing some free time activity with friends around."

The approach reported in this work uses Description Logics (DL) to describe context and preference. DL [Baader et al. 2003] is a decidable fragment of first order logic and is used in ontological languages such as OWL. A DL knowledge base consists of two components, a TBox, and an ABox. The TBox is essentially the vocabulary, and contains terminological knowledge about concepts (e.g. Child, Female) and roles (e.g. hasRoom, hasActivityType). The ABox contains assertions about individuals (e.g. Peter, Room5). Concepts and roles can be atomic or constructed using constructors $\cap$ (for intersection), $\cup$ (for union), and $\neg$ (for complement).

For example the context-based preference information "Peter prefers TV programs of human interest genre when doing some free time activity" is represented as a tuple (Context, Preference), where each of Context and Preference is a DL formula as follows.

Context: $\{\text{Peter}\} \cap (\exists\ \text{hasActivityType.FreeTimeActivity})$
Preference: $\text{TvProgram} \cap (\exists\ \text{hasgenre.}\{\text{Human-Interest}\})$.

The context-based preference information "if Peter is in the room with (at least) a friend then Peter prefers TV programs of a genre that is of common interest to him and the friend" is represented as

Context: $\{\text{Peter}\} \cap (\exists\ \text{hasFriend.}\ (\exists\ \text{hasRoom.}(\exists\ \text{roomOf.}\{\text{Peter}\}))\cap \text{v})$
Preference: $\text{TvProgram} \cap (\exists\ \text{hasgenre.}((\exists\ \text{tvInterstOf.}\{\text{Peter}\}) \cap (\exists\ \text{tvInterestOf.v})))$,

where v is a variable.

In the implementation, each concept is represented as a table, which is named after the concept. The table has one attribute, called ID, and is populated by individuals corresponding to that concept. Roles are represented similarly. For example, Tables I and II are representations of the two concepts, Person and hasActivityType.

Table I contains static information, and Table II contains dynamic information that must be received in real time from external sources. Using such tables, contexts can be expressed as SQL queries. If such a query returns a non-null answer, then the context is identified. The DL constructors $\cap$, $\cup$, and $\neg$ are mapped to SQL constructs INTERSECT, UNION, EXCEPT, and quantified sentences $\exists$R.D and $\forall$R.D are mapped onto nested SELECT statements.

The architecture takes the user query and information from the environment (via updated context tables such as Tables I and II) and determines the context. Then it uses the context information to select a preference. The preference is then used in two ways: to augment the user query (the pull mode) for final evaluation, and to generate triggers (the push mode) whereby the system proactively provides (pushes) information to the user. The first is used, for example, to augment the query with a specification of the genre of TV program when the user requires a TV program in a recognized context. The second is used for example, to trigger retrieval of the background information of the person who has just entered the room where the user is. The system is implemented on top of the DB2 database management system.

7.2. Planning

Planning concerns the problem of how to achieve a goal state starting from a known initial state. A plan is a sequence or partially ordered collection of actions that if executed starting from the initial state, is expected to achieve the goal state. There are many planning algorithms and techniques, for example Graphplan [Blum and Furst 1995], which analyzes graph-based search spaces of possible actions, hierarchical planning techniques [e.g. Tate et al. 2000], which predefine groups of goals and actions, and reactive planning [e.g. Firby 1999, Georgeff et al. 1985], which senses the environment and uses the information to adjust the plan.

There are several ways that plans and planning can be used in AmI scenarios, for example in Smart Homes. Planning can be used to coordinate the capabilities of the available resources to provide a solution or perform a task. Planning for AmI may have to deal with multiple agency; an example of this is provided in Amigoni et al. [2005]. Plans can be used, for example, to

—provide task guidance and reminders to inhabitants;
—allow AmI systems to share task execution with inhabitants;
—identify emergencies when the inhabitant is not doing what they are supposed to be doing or is not doing it correctly.

A discussion and survey of this is provided in Simpson et al. [2006]. This discussion is used liberally in the following overview.

An example of a system that uses plans to provide task guidance is COACH (Cognitive Orthosis for Assisting aCtivities at Home) [Mahilidis et al. 2001, 2003, 2004]. COACH does not do any planning, but it has a hand-coded internal representation of detailed steps taken to perform a task, for example the task of washing hands or making soup, stored in a plan library. It uses a video camera for locating the user's hands. This is used as input to a probabilistic neural network, which categorizes patterns of hand locations into steps. It then uses a simple algorithm to recognize plans from sequences of steps. Briefly, in the algorithm, the sequence of steps is compared to COACH's library of plans. If a match is not found, the most recently observed step is ignored and a new match is sought and so on, until a match is found. The system then guides the user through the remaining steps of the plan. As COACH does not actually perform any planning it cannot do any replanning when faced with unaccustomed circumstances. Also its plans are all sequences of actions, where the time of each step is relative to another, but it does not reason about actions that must occur at some absolute time.

The COACH system has been extended [Bouchard et al. 2008] to incorporate plan recognition for people with Alzheimer's disease. One complication in the case of such people is that when they are observed to do an unexpected action the action may be due to an interleaving of plans for different goals, or it may be due to an error on their part. To account for this, Bouchard et al. consider not only a plan library, but a space of hypotheses that includes compositions of pairs of plans in the library. The approach is based on probabilistic description logic [Heinsohn 1994], whereby a lattice structure of hypotheses is defined, augmented by probabilities that are prespecified from historical data.

PEAT (plan execution assistant and trainer) [Levinson 1997] is another plan-based system for providing guidance to the user. It is also capable of rescheduling activities in the light of unforeseen circumstances, but its information about the outside world comes only through the user; there are no sensor inputs.

ACHE (Adaptive Control of Home Environment) [Mozer 2005], a project at the University of Colorado, uses neural networks to monitor the relationship between time and date and operations of light switches and temperature controls. This is used to encode a sequence of actions within the neural network. This sequence is not recognizable as a classical plan, but can be used for monitoring the inhabitants' actions and for operating light and temperature switches. A cost evaluator guides ACHE's decisions by calculating an expected cost based on energy consumption and on user discomfort, interpreted as the likelihood of the user having to control lights or temperature manually.

Patkos et al. [2007] provide a brief survey of planning techniques with the aim of highlighting the shortcomings of existing techniques with regards to AmI applications. They identify several complications imposed on planning by AmI environments and applications. These include the need to generate and execute plans in a dynamic environment where the duration and success of action executions cannot be guaranteed. Also, in a multiagent situation, common to AmI environments, it may not be known a priori which agents can participate in the planning phase (an issue also discussed in the context of Amigone et al. [2005]) and which are available at the execution phase. Another challenge is whether or not the different agents are participating in planning share plan representation techniques. Patkos et al. conclude by suggesting two broad directions of research in this area. One concerns enhancements, along the lines of the continual planning approach [Desjardins et al. 1999], whereby planning and plan execution are combined. Here one plan is considered (as opposed to multiple plans for different contingencies), and continually updated as the execution and other information suggest opportunities as well as causes of failure. Another is context-aware planning, adding richer dynamic domain-specific factors to the more conventional factors, such as precedence constraints and plan length, to guide choices in planning.

Amigone et al. [2005] also discuss the challenges that arise when planning in an open ubiquitous environment. Primary among their concerns are that, (a) the devices available in the AmI environment may have limitations of processing and communication power, and (b) the number and type of devices connected to the system may not be known in advance and may change dynamically, for example when a person with a smart cell phone enters a room equipped with an AmI system, the cell phone can temporarily become part of the system. A classical planner typically iteratively decomposes the tasks until it reaches subtasks that are executable. In an AmI environment however, planners have the added complication of needing to keep track of the changing availability of devices and their capabilities, while decomposing tasks and building plans. To this end Amigone et al. propose a planner called Distributed Hierarchical task Network (D-HTN), which extends planning via hierarchical task networks (HTN).[13]

D-HTN is a centralized planner in that it centrally controls the plan building activity and it manages and takes into account distributed capabilities provided by the distributed devices, some of which may be permanent and some transient. So D-HTN incorporates awareness of the context within which it performs planning. The centralized planner and the devices are considered as cooperative agents. The architecture implemented is as shown in Figure 3.

The scenario considered is care of diabetic patients at home. Body monitoring devices provide data about the state of the patient, and this in turn, may require action such as adjusting the temperature of the room or contacting medical help or a specialist. The

[13]The best known HTN planner is probably Sacerdoti's NOAH [Sacerdoti 1977].

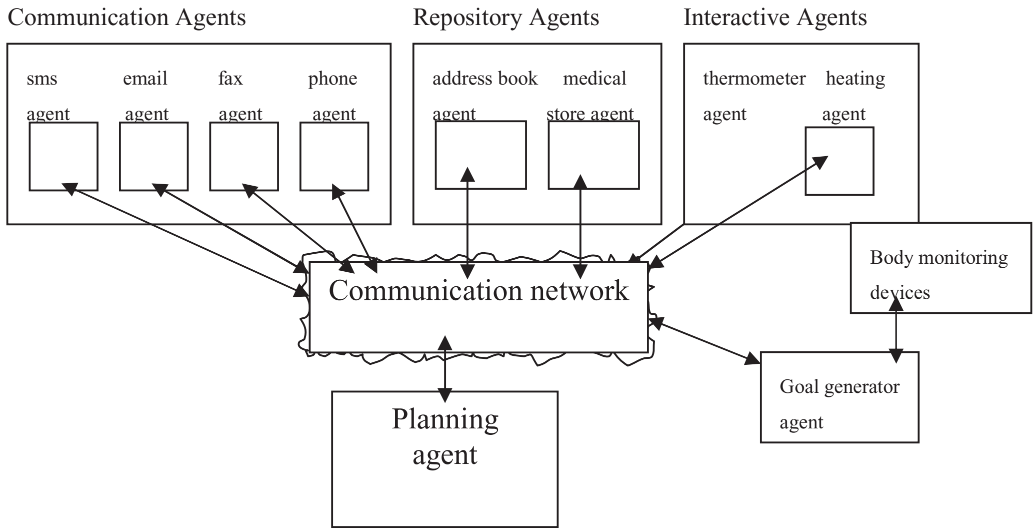

Fig. 3. An architecture for planning in an open ubiquitous environment.

Goal generator agent simulates any device or human that generates top-level goals for the Planning agent. The other agents are organized into three groups, the Communication agents, which can be transitory, the Repository agents representing the address book and a database of medicines available in the environment, and the Interactive agents providing sensors and actuators to check and adjust the room temperature.

HTN planners deal with task networks, which are represented as:
$[(n_1 : \alpha_1), (n_2 : \alpha_2), \ldots\ldots, (n_k : \alpha_k), \phi]$,

where each α_i is a task, primitive (executable) or otherwise (further decomposable), each n_i is a label identifying the task, and ϕ is a set of constraints on the tasks, for example $n_i < n_j$, to mean n_i has to be executed before n_j, and (n_i, p, n_j) to mean that p has to persist between n_i and n_j, that is p must be true after n_i and all the time until before n_j. An HTN planner starts with an initial task network consisting of the initial goal and a set M of available decompositions, each of the form m = (t,d), where m decomposes a task t to a task network d. M is in effect, a plan library. The planner repeatedly applies the decompositions to nonprimitive tasks until the resulting task network consists only of executable tasks.

In the architecture shown in Figure 3, the planning (D-HTN) is an extension of HTN, where the decomposition methods are distributed among several agents. The Planner agent is given a goal. It sends a message to all devices known to be currently part of the AmI system, asking for their decompositions of the goal. Awareness of which devices are available is implemented through JINI and its discovery mechanism. Any device that has the capability to form some decomposition of the goal may answer. Their answer will consist of a decomposition together with some quantitative measurements, such as effectiveness of plan, cost (resource consumption), and probability of success. The Planner then chooses from amongst, these, based on the quantitative values provided, and iterates if further decompositions are necessary, and finally allocates primitive tasks among the agents.

A brief discussion of some theoretical results is provided regarding decidability and complexity of the problem of whether a plan exists for a given goal, given a set of decompositions, an initial state, and a set of primitive tasks. In general the problem is shown to be semidecidable and exponential, and with certain restrictions decidable and polynomial.

A prototype implementation is described, with example goals such as “check and order insulin” and “call a technician to fix the insulin pump,” whereby the cooperation of the Repository and any available Communication agents is required. The paper concentrates on planning and not plan execution or plan repair. So for example, if after planning and before execution, an agent involved in the execution leaves the system, the execution aborts and a new planning process starts for the same top level goal.

7.3. Self Organization

As with Amigone et al. [2005], Encarnacao and Kirste [2005] are concerned with the dynamic nature of AmI environments where devices enter or leave. They argue that in the future, for AmI systems to be useful, they need to have the flexibility and adaptability afforded by self organization. A system that is carefully predesigned for anticipated scenarios cannot deal with unforeseen changes of devices or appliances in the environment. They argue that what is required is a system that is capable of self organization into ensembles constructed from the available resources, to be responsive to user needs and desires.

Some static predesigned AmI systems are already available or at least easily feasible. For example, in a smart conference room the projector may switch to the next speaker’s presentation as he approaches the front, dim the lights when the presentation starts and turn them up again when the presentation is finished. In such a scenario the devices work together according to an ensemble anticipated by the designer. This means that a change in the requirements or devices will probably have to be dealt with manually and possibly rather clumsily by the programmers. Self organization may be a solution for overcoming such shortcomings.

To enable self organization, the authors propose goal-based interactions, as opposed to the more customary function-based (or action-based) interactions such as “turn on,” “turn off,” “play,” and so on. With goal-based interactions, the user should be able simply to specify his goal, rather than to specify a sequence of actions that if executed would achieve his goal. A motivating example is for the user to specify the goal “I want to watch (the film) Chinatown now” and have the ensemble of devices work out which devices and what sequence of functionalities (actions) are needed. For example, turn on TV, turn on VCR, select video, position video at start, adjust air conditioning to a comfortable temperature, adjust sound level, adjust room lights, set TV channel for video viewing, and so on. In this way, goal-based interaction abstracts away from devices and functionalities, and allows the interaction with the system to be based on the user’s view/goals rather than on the system’s view of the world.

Goal-based interaction requires intention analysis, that is translating the user desires into concrete goals (possibly also taking context information into account), and strategy planning, mapping the goals to sequences of actions by devices. This latter can be realized through an explicit modelling of devices by precondition/effect rules, which turns the strategy planning problem into a familiar, classic planning problem.

Based on these concepts, the authors have produced an experimental system called SODAPOP (Self-Organisating Data-flow Architectures suPporting Ontology-based problem decomPosition). The work has been done within a German-based project called DYNAMITE (Dynamic Adaptive Multimodal IT-Ensembles). SODAPOP is a middleware for self-oganizing AmI systems. It consists of channels and transducers. Briefly, the channels are message busses and the transducers are the devices. The transducers subscribe to channels by declaring what messages they are able to process and whether they can allow other transducers to process the same message concurrently. When a channel receives a message it decomposes it via its strategy planning to simpler messages that the subscribing transducers can handle and sends these to them. The channels are capable of receiving and decomposing goal-based messages and the publish/subscribe approach of the transducers allows on-the-fly formation of ensembles. A system example is given containing a TV, a stereo, a speech input device, a display unit, and an avatar. With this setup, for example, the audio of a film may be rendered through the stereo if its audio quality is judged superior to the TV. Also all outputs may be rendered through the avatar.

7.4. Embodied Systems - PEIS

Saffiotti and Broxvall [2005] report work in combining robotics, ubiquitous computing and AI technologies, to provide not just a connected environment of interacting devices, but with emphasis on performance of physical tasks by robots. The intention is to provide Physically Embedded Intelligent Systems (PEIS), in what they call a PEIS ecology, namely a cooperative physical environment facilitated by communication.

A motivating scenario is given in the following. Johanna is 76 years old and lives in a small house. Just before she wakes up, her fridge realizes that there is little milk left. It sends a request for a bottle of milk to the local store and Johanna's autonomous trolley goes to the store to collect it together with the usual daily newspaper. When Johanna gets out of bed her motion is detected and the coffee machine starts to make coffee. A team of robots brings her breakfast items to the table. As Johanna leaves the bathroom the cleaning robot starts cleaning it.

To study the viability of such scenarios and concepts, the authors have constructed a PEIS-Home, which is a mini-apartment with a living are, kitchen, and bathroom. The ceilings have been lowered in order to provide space above them for cables and computing equipment, and to provide observation spaces. The Home is equipped with two Magellan Pro robots, a tracking component and a monitoring component. The tracking component is a localization system using Web cameras mounted on the ceiling to track colored objects; the robots have colored tops. The monitoring system is for the use of the human experimenters to observe the activities inside the Home, and it consists of a visualization tool and a record/playback tool. Both robots, Pippi and Emil, have communication and fuzzy logic-based negotiation functionalities. Emil, in addition, has a conditional planning functionality.

In one reported experiment, Emil receives a request to wake Johanna up. It delegates the task to the other robot, and generates a plan for the waking up consisting of three steps, go to bed, talk to Johanna, go to sofa. It tells Pippi of the first step. Pippi executes the action, keeping track of its whereabouts by the tracking device, and notifies Emil when it is at the bed. Emil communicates the next action of the plan to Pippi, and so on, until the task is accomplished. The communication between the devices assumes a shared (tuple and event based) ontology.

7.5. Decision Trees

Stankovski and Trnkoczy [2006] discuss the use of decision trees in smart home applications. The main uses identified are in recognition of usual and unusual events in smart homes and prediction of next events.

Decision trees provide a graphic representation of possible combinations of attribute values (Figure 4 is a simple example). They may be user friendly because of the visual nature of the representation. Furthermore, it is straightforward to translate them into rules for automated reasoning. A decision tree is a collection of tests. The nodes represent tests of given attributes and the branching at each node represents the possible values of the attributes. The leaves provide classifications of the instances.

In smart homes, sensor data can be collected regarding the whereabouts of the person, his interaction with appliances (turned light/cooker on/off) and the duration

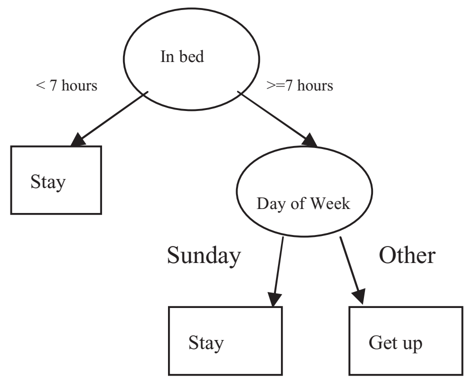

Fig. 4. A decision tree.

of events (light on for 4 hours, cooker off for 3 hours). This data can then provide a training set for creating decision trees, a process called induction of decision trees (a well-known algorithm for induction is the TDIDT/ID3 [Mitchell 1997]). The resulting decision tree can then be subjected to postpruning either simply visually or using a standard decision tree pruning algorithm.

The assumption here is that the decision tree and the set of values used for its induction represent normal activity. So in the example shown in Figure 4, it might be expected that if the person stays in bed for more than 7 hours on a Monday, it is recognized as unusual behavior. Decision trees are appropriate for domains that can be represented by attribute-value pairs, with finite value sets and a predetermined set of attributes. Also they work best if the set of possible values is small; many algorithms support only binary splits. Furthermore, incorporation of an explicit notion of time remains a research issue. The paper does not report any concrete application, implementation, evaluation, or experimentation.

8. AMI: HUMAN INSPIRED AND AFFECTIVE COMPUTING

Affective computing is concerned with enabling computers to recognize human emotion and to respond accordingly. Early work on affective computing can be found in Salovey and Mayer [1990]. Emotionally intelligent computers have clear advantages and applications in AmI. Picard [2000, 2007], describes the use of physiological signals and explicit user self-reports to help computers recognize emotion. The same papers also report how different types of human-computer interaction may affect human activity, for example people may interact for longer periods with the computer if it appears to understand and empathize with their emotions. The European project (FP6) e-Sense (Capturing Ambient intelligence for Mobile Communications through Wireless sensor Networks,[14] and the later (FP7) Sensei (Integrating the Physical with the Digital World, of the Network of the Future,[15] which started in 2008, are partly concerned with remote emoting, aiming at measuring emotions, and transmitting them to others, so that, for example, friends and family can share the excitement of each other's sport and leisure activities.

Forest et al. [2006] describe context-awareness afforded by wireless sensor networks as awareness of a collection of factors, including emotions. They outline several scenarios for mood-based services, where the mood is determined via physiological factors,

[14] http://www.ist-esense.org.
[15] http://www.sensi-project.eu/.

such as ECG, breathing rate, and alterations in voice and in skin conductivity. In one scenario, feelings of distress and fear are recognized and lead to an assumption of physical danger, which in turn, triggers automatic notification to an appropriate person, for example a neighbour or the police. Another scenario, commercially popular, is in games, capturing the emotions of the player and amplifying it within the agents in the game, thus enhancing the player's experience and feeling of immersion. Another scenario involves social networking, where people can have body area networks (BANs), which include body sensors that monitor the physiological factors and classify the users' moods. If BANs are also equipped with information about their users' preferences (for example for film genres, etc.) the BANs can interact with one another in public situations and find matches for their human users.

Gonzales et al. [2006] outline wearIT@work, an EU funded FP6 project,[16] which explores the use of wearable computing devices in emergency responses. A particular interest in the project is how such devices can help determine the emotional state of firefighters. This information is then combined with the profiles of the wearers to help the commanding officers manage the deployment of the firefighters in emergencies. Other groups [for example Marreiros et al. 2007a, 2007b] have been looking into how emotional factors can be incorporated in a system that simulates ubiquitous group decision-making.

Within HCI, various techniques have been used to recognize human emotions. Some have concentrated on facial expressions and speech [Cowie et al. 2001], and others on physiological parameters such as heart rate and electrodermal activity [Kim et al. 2004]. There are two prominent approaches to emotional modelling, the Basic Emotions approach [that is Oatley and Johnson-Laird [1987] and Ekman [1992] and the Dimensional Approach [Vastfjall et al. 2000]. The first assumes there are several (six has been the more widely accepted number) basic emotions that can be distinguished from one another, such as anger, happiness, sadness, fear, disgust, and surprise. The second assumes the existence of several (normally two or three) dimensions that are sufficient for characterizing and distinguishing different emotions. The dimensions include valence with positive and negative poles (indicating pleasantness versus unpleasantness), arousal with calm and excited poles, and dominance with weak and strong poles (indicating vulnerability versus dominance). For example, both anger and fear are considered to have negative valence and high arousal, but anger has strong dominance and fear has weak dominance.

8.1. Herbon et al. [2006]

This paper describes experimental work that compares the two emotional modelling approaches mentioned previously, namely the Basic Emotions approach and the Dimensional Approach. The paper also explores the correlation between physiological parameters and self-reports. The experiments involved showing subjects 5 films, each about 4.5 minutes long, intended to arouse emotions. For example they showed a short funny cartoon to induce a positive arousing emotion (like happiness). The subjects (40 participants at the University of Surrey) were invited to self-assess after seeing each film. The self-assessment asked for both a rating on the dimensional scale and a choice of the basic emotions. The physiological parameters measured were ECG (electrocardiogram, breathing rate), EDA (electrodermal activity, skin temperature), EMG (electromyogram, facial muscle activity) and voice.

The two models of emotion were compared by asking the subjects to fill in a questionnaire expressing their preferences between the two. The results showed a slight preference for the dimension model over the basic emotion model in terms of ease of

[16]http://www.wearitatwork.com/Home.20.0.html.

use. This appeared to be due to the range of values available along the dimensions, enabling a more accurate rating of the emotion.

Results on the analysis of the physiological factors and emotion dimensions showed a correlation between increasing valence and facial cheek muscles contracting and pulse decreasing. Another correlation identified was between increasing arousal and decreasing skin resistance and increasing respiratory activity. But no correlation was demonstrated between brow muscles or other ECG parameters and either arousal or valence.

8.2. Picard [2000, 2007], Paiva et al. [2007]

These papers describe the work of the Affective Computing Research Group of the MIT Media Laboratory aimed at giving computers the ability to recognize and respond to human emotions. The emotions of interest in their work include frustration, anger, distress, joy, interest, and confusion. The papers describe how emotions can be determined, how user frustration may be dealt with once it is recognized, and some applications of affective computing.

Two approaches are described for determining user emotions: self-report and concurrent expression. In the first, the user stops what they are doing and actively communicates an emotion, for example a coffee machine which has thumbs-up/thumbs-down buttons via which the users indicate their satisfaction or lack of it with its operations. The feedback is recorded and provides information for the designers of the machine [Norwood 2000].

The concurrent expression approach detects emotion without the user having to stop what they are doing to report it. As an example, the group has developed wearable "expression glasses," which have a small point of contact with the brow and can sense changes in facial muscles such as furrowing the brow in confusion or interest. Another device described, is the Touchphone. This is a normal phone equipped with a pressure sensor, measuring how tightly a person is holding the phone. This measure is mapped onto a color, changing from blue to red as the pressure increases, which is transmitted to the other person on the phone. So the Touchphone may reveal some information about changes in a person's emotional or physical state, without actually revealing the state itself.

An experiment is reported on how human frustration is dealt with once it is recognized. The experiment involved 70 subjects who were asked to play a computer game to show their intelligence and also to win monetary prizes. Some of the subjects were then faced with particularly frustrating situations while playing the game, for example with network delays. Afterwards subjects would interact with an automated agent, and then they would go back to the game. They would be given a choice of either ending the game or continuing to play, in which case the length of time they continued with the game was measured.

There were three types of automated agents: ignore agent, which just asked about the game, ignoring emotions, vent agent, which asked about the game but also asked about the person's emotional state, giving room to vent, but with no active empathy/sympathy, and emotion support agent, which was similar to the vent agent, but showed active empathy ("sorry to hear your experience was not better," "it sounds like you felt fairly frustrated playing this game. Is that about right?"). The results of the experiment showed that people interacting with the emotion support agent played significantly longer after the interaction than those who interacted with the other two agents. The emotion support agent was not particularly intelligent or sophisticated, so it was interesting that such a relatively simple device seemed to ease human frustration.

Several applications of affective computing are described. One is a wearable DJ, which measures the user's skin conductivity while he listens to songs. Higher conductivity is associated with being energized and lower conductivity with being calmed. The data is then used to choose songs for the user to listen to when he asks for "some jazz to pop me up," for example. Another is a music conductor's jacket which through EMG (electromyogram, a test that is used to record the electrical activity of muscles) and respiration sensors maps patterns of muscle tension and breathing into features of the music. One proposed application of the jacket is for teaching student conductors, for example to give them feedback on their timing and tension.

8.3. D'Mello et al. [2008] AutoTutor

This paper describes work on an affect-sensitive tutor system, AutoTutor. AutoTutor helps students learn physics, computer skills, and critical thinking. It has a two-phase strategy, first recognizing affect and then responding to it. The affects considered, are boredom, frustration, confusion, and attentiveness (which they call "flow"). These were chosen as more relevant to learning than the standard basic emotions [Oatley and Johnson-Laird 1987; Ekman 1992] which are anger, fear, sadness, happiness, disgust, and surprise.

Affect detection uses three sources of information, conversational cues, body language, and facial features. For conversational cues, certain features are monitored, such as time taken by the student before giving a response (temporal feature) and the number of characters or speech acts used (response verbosity). For body language information, pressures on the seat and back of the chair of the student are monitored, with the assumption that increasing pressure on the seat signifies attentiveness (student is positioning his body closer to the screen), and increased pressure on the back of the seat signifies low attentiveness (student is leaning back and further away from the screen). For facial features, the system tracks the pupils of the eyes, fits templates to the eyes and brows, and then labels facial action units. The assumption is that confusion is associated with a lowering of the brow and tightening of the eyelids. How to combine the information from these three sources and resolve any contradictions, for a mapping to an affect, is a subject of the group's ongoing work. This is also a research topic within other projects, for example e-Sense [Forest et al. 2006].

AutoTutor's response when an affective state is detected is specified by production rules. Production rules are similar to ECA rules. They are of the form If Conditions then Actions, and if Conditions are matched with attributes holding in a state, the production rule is triggered, that is becomes a candidate for execution. Some selection strategy will then determine which of the triggered rules will actually fire, meaning that their Actions are executed.

In AutoTutor the contents of the production rules have been inspired by two theories of learning; attribution theory [e.g. Batson et al. 1995] and cognitive disequilibrium theory [e.g. Craig et al. 2004] and by feedback collected through experiments with pedagogical experts. The conditions of the production rules refer to the student's performance and detected emotional state, and the actions refer to primitive actions within the capabilities of AutoTutor, which include speech intonation (e.g. pitch, intensity), facial expression (e.g. surprise, delight) of the AutoTutor's embodied avatar, and a dialogue move (e.g. hint, prompt). An informal example of a production rule is the following.

> If the student has been doing well overall, but not in his most recent responses, and the current emotion is classified as boredom then AutoTutor might say: "Maybe this topic is getting old. I'll help you finish so we can try something new."

AutoTutor's phrases in responding are chosen randomly from a list of similar responses to these conditions.

As well as affective computing, there is work within AmI on architectures inspired by theories of human consciousness. Two examples are Piva et al. [2005] and Marchesotti et al. [2005], both inspired by the model of human consciousness given in Damasio [2000]. This human model describes two broad interacting components, one that is nonconscious, primarily focused on maintaining the body within a narrow set of parameters required for survival, and one that is conscious, that deals with memories of the past and anticipations of the future.

Marchesotti et al. report an implemented architecture for complex event classification based on Damasio's model. The proposed artificial architecture contains two sets of sensors, one corresponding to the human nonconscious component, for observations of the internal parts of the organism and one corresponding to the human conscious component, for observing external parts. The organism in the implementation is a lab at the University of Genova. The first set of sensors in this context corresponds to software sensors monitoring the number of machines switched on, and the network load. The second set of sensors corresponds to cameras covering the lab and its two access points and sensors monitoring keyboard and mouse activity. The objective behind the architecture is the fusion of the sensor data to recognize events such as people entering the lab and low or high human and machine activity.

Piva et al. also use a university lab as a test bed for their implemented architecture. They also employ two sets of sensors, one monitoring the state of the lab and the machines, the other monitoring human movements and human queries. The objective is to interact with the humans, proactively or in reaction to their queries, for example about availability of machines in the lab, using different modes of interfaces. The interaction is based on rules similar to production rules. The choice of interface is context-dependent. For example if the query is about whether or not machines are available for use in the lab, the response can be delivered via the user PDA if the user is away from the lab, or via a display unit at the entrance of the lab if the user is at the lab.

9. SOCIAL, ECONOMIC, AND ETHICAL ISSUES

Ambient intelligence involves extensive and invisible integration of computer technologies in people's everyday lives. Such integration will inevitably open up issues of privacy, risk, acceptance, and security. We have already seen some of the negative implications of internet technologies in cyber crime, pornography, including child pornography, and spam email nuisance. But, given the pervasive nature of AmI, its social and ethical implications go much further. It has been widely acknowledged that there is a need for acceptable standards and for laws regulating access, to avoid social and ethical problems.

RFID tags alone, have already been the subject of much debate and controversy. RFID (radio frequency identification) is a tracking technology that uses computer chips, which can be embedded into, or fixed to, more or less any physical item. RFID tags can be very small, the size of postage stamps or smaller. Those without an independent power source can transmit information from a few inches to a few feet away, and those with attached batteries can transmit information further. Many aspects of RFID tags and associated technology alarm some writers [e.g. Albrecht and McIntyre 2006; CASPIAN]. For example, unlike bar codes, proposed RFID tags can assign unique IDs to each individual item, the purchaser of which can be determined from store loyalty card information. Tags can be readable from a distance and through material such as wallets and clothing. Moreover, there are the unknown consequences of exposure to the electromagnetic energy emitted from large quantities of such tags.

Wright et al. [2008] discuss the security and privacy threats of ambient environments and technologies and illustrate these by four "dark scenarios" (SWAMI dark scenarios - Safeguards in a World of AMbient Intelligence). Among the threats illustrated in these scenarios is the criminal usage of AmI, for example someone's house security systems being disabled by criminals for the purposes of blackmail and extortion, or several employees of a major data mining company pooling together their partial security clearances allowing them to steal personal records of millions of people spread over several countries. Another example involves burglars obtaining details about a woman, including her home address and the facts that she lives alone and is well-off, from the RFID tag on a blouse she buys combined with some details on a database system they hack into.

Some scenarios in Wright et al. [2008] deal with issues posed by AmI in travel and healthcare. These concern overreliance on AmI, lack of interoperability, and disproportionate requests for personal information. For example, an accident involves a busfull of people touring a city. The accident itself, is later found to be due to a security lapse, allowing some kids to break into the city's traffic management system. The personal health monitoring devices of the people caught up in the accident help the paramedics to identify those in most urgent need of assistance. But there are people whose health monitoring devices are not up-to-date models, and this leads to fatal misdiagnosis by the paramedics, who are all too reliant on the AmI devices and no longer exercise the basic skills of diagnosis. There are also tourists whose devices are not compatible with the city's systems. Once the patients reach the hospital, they are asked to provide access to all their personal details, or to sign a statement relieving the hospital of any liability regarding their treatment.

Kemppainen et al. [2007] concentrate on the potential of AmI for people with disabilities, for example in telemonitoring and care through smart homes. The authors call for guidelines in the design of systems and interfaces that would ensure such minority groups are not excluded from the benefits of AmI. They draw attention to the fine balance needed in cases of people with disabilities, in the sharing of control between the individual and the system.

In the following, we survey some articles that address the issues of social and privacy impacts of AmI technology. The first two, Bohn et al. [2004] and Friedewald et al. [2005], identify the two central features of AmI that pose the main challenge to privacy. These are ability of AmI systems to collect large and detailed amounts of data about individuals' everyday activities over long periods of time, and the enhanced ability for integrating, searching, and retrieving these large amounts of data. These features are central for one of the key objectives of AmI, which is to provide personalized services. Another article, Van Heerde et al [2006], provides a small implemented case study of how some of these concerns may be addressed, and another, EU project e-Sense, provides focus group studies of user attitudes to such concerns. The final article, Rouvroy [2008], looks at difficulties in law in dealing with AmI issues, and points to a need for interdisciplinary (including law and information technology) discussions.

9.1. Bohn et al. [2004]

In this interesting and very readable article Bohn et al. examine the technological and economic driving forces behind AmI and discuss its potential implications from social, ethical, and economic points of view. They identify several advancing technologies as the primary technical driving forces behind AmI. These are advances in microsystems and nanotechnologies, providing very small and cheap processors with integrated sensors and wireless communication capabilities, and electronic markers (such as future descendants of smart labels and RFID tags, already used as antitheft devices in department stores and libraries). Such markers would allow real-time links to information sources and databases held remotely, ultimately allowing identification of individual objects and their association with a wealth of data.

They imagine one way of implementing such “communicating objects” by equipping them with electronic labels containing Internet addresses as digital information. Then one might read the information at the Internet address by pointing a portable device at the object, giving the impression that the object itself is transmitting the information. The information may be operating instructions or medical advice. It could be personalized and context-dependent, depending on the characteristics of the person with the portable reading device, their age, and preferred language. But the authors conjecture that the information may also depend on whether or not they have paid their taxes. Such possibilities have led to concerns about these technologies [Thackara 2001].

They imagine what an “ambient economy” may be like, resulting from the information that can be associated with all objects and the pervasive and active nature of devices. Entities in the economic process, for example goods or vehicles, can be tracked throughout their lifetimes, from the start time of their components onwards. Smart objects can be provided with self-conscious capabilities allowing them, for example, to monitor their storage environments for parameters such as humidity and temperature. Their communication capabilities may further allow them to alert the outside world if the environment becomes unfavorable. Going even further, they may have the capability of actively altering their environment. Such capabilities can significantly reduce wastage, and the cost and time needed for replacing damaged goods.

Smart goods can provide a history of themselves, used cars can provide information about replacement and usage of their parts, organic foods can provide information about their source and cultivation, thus, provided the information is reliable, allowing the consumer to buy with confidence.

AmI could bring about new business models taking us beyond the current e-business model. Objects could subtly advertise themselves. As you walk in the street and you see someone wearing something, you could get information about where you can buy it and how much it costs, from the invisible tags embedded in it. Maybe the wearer can even earn a commission when someone nearby sees it and buys it. There is the possibility of cross-advertising in the ambient environment, an electronic version of what already happens. For example, currently some washing machine documentations suggest using a particular water softener. But as embedded tags and devices could provide access to a wealth of background information, this cross advertising could go much further. A smart refrigerator could recommend recipes based on what is available in the fridge and what is discounted at the nearby, or the fridge’s chosen, supermarket. Prices of everyday goods could be tailored to individual buyers based on profile information available on them, in turn determining how much they are prepared to pay.

Consumer profiles can provide a useful service by screening information that is unlikely to be of interest. But another view of profiles is that they, in effect, help other parties decide what information to withhold from individuals. There is also the obvious risk of accidental or intentional leaking of the profile and data used to compile it. The ultimate in shopping, according to the authors, can be when objects do their own shopping, “autonomous purchasing objects” in “silent commerce,” according to the business consultancy Accenture. For example, photocopiers could order their own paper supplies, but then possibly also Barbie dolls could order their clothes.

Pay-as-you-go or pay-per-use is already a model in use, for example with phones and public utilities. The authors surmize that AmI may extend its usage. A sofa or a bed, for example, may be able to monitor its usage, frequency, and weight of people using it, and issue periodic bills according to usage. This model may be attractive to corporate customers such as hotels and offices, possibly less to domestic customers. The possibility of monitoring and recording activity can bring other uses, and associated risks to privacy. For example a smart car can be manufactured to provide detailed records of its owners driving habits and parking style, thus allowing insurers to use much more detailed criteria than age or location for setting premiums. In the case of a claim the cars can provide information about the circumstances of the accident, thus affecting the payout.

AmI can provide sophisticated support for everyday living, but the information capabilities it may use for this purpose can also potentially provide an invisible and comprehensive surveillance network—walls literally can have ears. This surveillance will not be particularly for unusual activity but for ordinary regular everyday activities. The authors argue that given the potential economic advantages to the sellers, it may be made difficult to opt out of the AmI environment, to the extent that opting out becomes a luxury for those who can afford it, for example having to pay higher insurance premiums for cars with less pervasive AmI capabilities.

The authors go on to identify three additional issues.

Reliability. AmI landscapes are by definition embedded and unobtrusive. This could make it difficult to notice that some part of it may not be working as required. The analogy they give is with a telephone. If you pick a phone up and there is no dial tone you know it is not working, but detection and diagnosis of faults may be much more complicated with invisible networked devices.

Delegation of control. Who is ultimately in control of the contents of AmI devices? Who is legally liable for the fridge that orders unwanted food, or the toy that influences the child by interacting with him?

Social compatibility and acceptance. How error prone will these environments be? What is the impact on health and environment of such large collections of tiny, individually harmless, electronic tags surrounding us, for example in the supermarket products and ultimately in our garbage?

9.2. Friedewald et al. [2005]

This work involves a study of more than 70 Research and Development (R&D), mostly EU-funded, projects and roadmaps, from the point of view of what types of scenario they focus on, what assumptions they make about the users, and the control of the AmI systems they envisage. The projects cover five application domains: home; health; shopping; work and mobility; leisure and entertainment, and focus on different levels of detail, from detailed scenarios, much like screenplays, to broad application scenarios, to even broader functionalities and trends.

For each application domain, based on the projects they have studied, Friedewald et al. identify a collection of services the AmI support is aimed at providing. For example in the last, leisure and entertainment, the aims could include:

—enhancing and personalizing the experience of visiting museums and other sites [e.g. Chervest et al. 2002];
—providing context-aware entertainment—for example selecting music that fits a person's mood, using quantifiable musical features, such as tempo [e.g. Sleeth 2002].

In the work and mobility domain the aims could include [e.g. Heinonen 2004]:

—turning the whole office space into a smart environment capable of tracking people and recording their work;
—increasing safety and security,
—to support workers' mobility, allowing them to work from anywhere and anytime.

One conclusion of the analysis provided in this study is that most applications assume a single user who has accepted, and can afford, AmI technologies. Even in the home environment, account is not taken much of families being present, and members having different preferences. Another conclusion is that the implementation of AmI features in work environments is in a more advanced state than in home environments. The study also looked at the level of personal (human) control over AmI systems. It concluded that in the envisaged and developing applications, where the AmI system was aimed at providing safety or security it had a high level of control. In particular, AmI control is assumed to be very high in envisaged emergency situations, requiring little communication with humans. On the other hand, where the system had a more advice-giving role it had lower levels of control, possibly subordinate to the user.

Other features of AmI studied in this paper concern information flow, its advantages and dangers. There is already a steady one-way flow of information from consumers to sellers, where sellers get access to, and store, sensitive data about customers. Businesses collect personal data about people's shopping, travelling, Web surfing, and TV watching behavior in order to compile profiles. AmI massively increases the amounts of detailed personalized data that is collected and stored, and has the potential to make, and indeed in some applications must make, such data easily available. For example in the case of a medical emergency, sensitive medical and identification information can be sent via local and wearable devices. Thus AmI may end up facilitating invasion of privacy and identity theft, because of the increasing amounts of personalized information stored, and its means of transmission via personal devices.

Location tracking systems, for example, help people find a near by acquaintance, but in the wrong hands, it can lead to attack or kidnapping. AmI gives parents a great deal of control and surveillance over their children. But at what age should this cease and who decides? An intelligent bed monitoring the weight of the person sleeping on it can be used to monitor the weight of elderly people to help avoid damaging weight loss, but they can also be used to spy on people, for example to detect how many people are in the bed. Furthermore, as Bohn et al. [2004] have also observed, personalization of data and provision of services, can ultimately lead to the control and filtering of what news or information the users see.

Dependence on AmI in general, can lead to loss of the ability to manage one's life. Assisted living and automated healthcare for the elderly, in particular, despite their obvious benefits, can lead to such loss, and telecare can contribute to feelings of isolation. Finally, the authors view the possible AmI divide in the opposite direction from Bohn et al. [2004]. Not that AmI will be for the masses, with only the rich being able to evade it, but that the high costs of AmI support will mean that only those who can afford it may benefit from it.

9.3. Van Heerde et al. [2006]

In this brief paper the authors look at data privacy in ambient intelligence settings. As with the other authors, they note that it is the high quality and large quantities of data that can be collected that enable the intelligence of AmI systems, and at the same time provide the privacy challenges. Moreover, the collection of this data through ubiquitous systems leads to asymmetric information [Jiang et al. 2002], where there is significant asymmetry between the donor and collector the proposal of Van Heerde et al. is to give information donors (those about whom information is collected) control over how long, and in what form, information about them is kept.

Their implemented case study concerns an organization (a university) collecting information about its members accessing Web sites. The purpose of this data is to enable ranking of Web sites and making recommendations to those with similar interests. However, the malicious use of the data can disclose information about what times, and for how long, someone accessed some given Web sites. To counter this, the proposed solution is that users can specify life-cycle policies on data collected on them. Life-cycle policies can ask for periodic replacement of the data with increasingly more abstract and less detailed versions. For example, a life-cycle policy can be that initially the data collected includes time breakdown and user ID and URL of the site visited. After one hour the data is downgraded by replacing the first two components to the hour in which the URL was visited, and the group to which the user belonged. After one month the group is replaced by the university, and the URL is replaced by a general classification of the Web site, and after one year the data is deleted.

9.4. EU project e-Sense

e-Sense (Full title: Capturing Ambient intelligence for Mobile Communications through Wireless sensor Networks)[17] is an EU-funded (Framework 6) project, consisting of a large (IP- Integrated Project) consortium of 23 partners, from industry (e.g., IBM, Fujitsu UK), management, SMEs (small or medium size enterprises), research institutes, and academic institutions (e.g. University of Surrey, TUDelft, Kings College London, University of Twente). Its objectives are to develop communication technologies using wireless networks, aimed at transmitting context data, such as biometrics information, for example heart rate, skin conductance, breathing rate, and data about the emotional state of the sensed participants.

The project included a study of the public reaction to such technologies and their acceptability. They used focus group techniques, much like Niemela et al. [2007] (Section 3.3 in this article), using audio-visual demonstrations followed by discussions. The audio-visual demonstrations can be found online.[18] They consist primarily of picture boards with headings. The study was conducted in three countries: France, Germany, and Spain, and on three applications, personal, healthcare, and industrial. For these applications, the focus groups consisted of the general public, healthcare professionals (doctors and nurses), and industrial sector professionals (transport, retailing, food processing, logistics).

For the personal applications, the audio-visual demonstrations focused on fictional systems that continuously monitored the user's conditions (mood, location, activity, social situation ...) in order to provide personalized advice and services, for example in entertainment and counselling. The study found that the general view of the focus groups was negative, finding the continuous monitoring too intrusive, and only possibly acceptable in cases of special needs. There were only minor variations based on nationalities.

For the healthcare applications, the demonstrations focused on scenarios involving health monitoring at home via body sensors, and involving a wireless ward, whereby nurses could keep an eye on patients via wireless monitoring. Here the response from the health professionals was positive, but with concrete and serious concerns, such as the legal liabilities of the doctors and the hospital, fear of decreased communication with the patients, and concern about the responsibilities for diagnosis.

For the industrial applications, the demonstrations focused on a range of scenarios, for example, the use of RFID tags for tracking and monitoring goods for logistics purposes, and in the shops for customer use for traceability and product information. Other scenarios included using sensor data in factories to monitor the quality and state of both equipment and products. Here the response of the professionals in all sectors that took part was positive, in terms of the possible enhancement of productivity and competitiveness, but less optimistic in terms of short term investment in such technologies. It was also felt that in the medium term (3–5 years), the focus of such technologies for industrial applications should be on production and transport, and in general for the benefit of the organizations and enterprizes (rather than the retail and customer side).

[17](http://www.ist-esense.org/index.php?id=18.

[18]http://www.msh-alpes.prd.fr/e-sense/.

9.5. Rouvroy [2008]

This paper considers the current European privacy and data protection frameworks and questions if they are applicable and adequate for dealing with the kind of data collection and processing that is at the heart of AmI scenarios and technologies. It focuses on two aspects related to personal data, namely “freedom from unreasonable constraints (from the state or from others) on the construction of one’s identity,” and “control over aspects of the identity one projects to the world.”

The article argues that the European human rights framework incorporates “autonomy in the construction of one’s identity,” explicitly in the right to privacy. One consequence of this, interpreted in courts, is the individual’s right to control personal information. The pervasiveness of AmI and the invisibility of data collection and information systems may make it highly unlikely that the individual (the person being observed) will retain control over the data. Moreover, the envisaged seamlessness and invisibility of AmI systems may also have the effect that they disappear from human consciousness, and in that way they bypass intentionality and control.

Furthermore, one objective of AmI systems is to learn user profiles in order to respond to human needs. But the paper argues, that these needs are being defined increasingly by the systems themselves, and thus by the designers of the systems, and not by the users. Moreover, the purpose of the learning is often to classify people into categories, which in turn, can affect the accessibility of information and services to people. Thus the paper argues that the AmI challenge to privacy and data protection is not just in the collection of data, but also in the interpretation of the data, which can constrain the user’s environment, choices and preferences.

Rouvroy shares with Bohn et al. [2004] the concern about delegation of control. AmI systems are likely to be distributed systems in which multiple artificial and human agents collaborate and interact. So the notion of human agency, traditional in law for assigning individual responsibility and liability, becomes blurred. In such interconnected systems the identification of where the blame lies when something goes wrong becomes a difficult issue.

A European directive identifies personal data as “any information relating to an identified . . . person.” Even though the concept of personal data is not unanimously agreed upon in all European Community countries, the European data protection framework identifies a category of sensitive data, and another directive makes it illegal to process such sensitive personal data, such as data revealing racial or ethnic origin and political opinion, and data concerning health or sex life. This is another aspect for discussion in relation to the nature of data collected in AmI scenarios. Images of a person can convey racial and ethnic origins, their preferred entertainment, Web sites and TV programs may convey their political opinion and religious tendencies, and what they buy may convey information about their health. In particular, the processing and integration of all such data about an individual can lead to the extraction of much more, possibly sensitive, personal data.

10. CONCLUSION

In this selective survey we looked at several application areas of ambient intelligence, including the smart home, care of the elderly, healthcare, business and commerce, and leisure and tourism. In these, we looked broadly at trends, requirements, and challenges, as well as at technical developments and implemented demonstrators. We also looked at studies of the attitudes of human target groups towards these applications and technologies.

Furthermore, in the context of AmI we looked at several data management and artificial intelligence technologies, including event-condition-action rules, production rules, learning, fuzzy logics, planning, plan recognition, temporal reasoning, and case-based reasoning. We looked at how current technologies are being used and what extensions are thought to be necessary. We also looked at several approaches to using agents, for example as abstraction tools, for modelling devices and their interactions, and as middleware.

We then considered the role of affective computing and human emotions in ambient intelligence. We considered different approaches to recognizing and classifying emotions, including self-reports, physiological metrics, seat and hand pressure sensors, and characteristics of speech acts. We also looked at studies correlating some of these different techniques, and analyzing people's preferences in rating their emotions according to the two main models of human emotions, the Basic Emotions and the Dimensional Emotional models. Furthermore, we looked at studies of how human emotions may be influenced by the way computer systems interact with humans. We completed the survey by exploring the social and ethical implications, and challenges of, ambient intelligence technologies.

There are two broad schools of thought regarding AmI. One is that much of the envisaged functionality is realizable through advances in hardware and sensor technologies, functioning with simple data and simple reasoning mechanisms. The other is that the full potential of AmI cannot be realized without sophisticated knowledge representation and reasoning and other AI and agent-oriented technologies. This survey has been biased towards the second school of thought. It has explored what AI and agent technologies can offer in processing, and in making decisions, on the basis of the data provided by the hardware.

Several concluding observations can be made from this survey. One, not surprising, is the universal agreement on the need for context-sensitivity in AmI systems. AutoTutor [D'Mello et al. 2008], for example, uses the context of the pupil's emotional state to decide what to do next. The planner of Amigone et al. [2005] constructs plans in the context of the currently available devices and their capabilities. All of the smart home and elder care systems we looked at decide what action(s) to perform in the context of the current circumstances, be it to adjust lights or heaters, provide advice about execution of a task, or to suggest a new schedule of activities to compensate for disturbances in previous schedules.

What is more surprising, or at least more interesting, is the variety of different techniques proposed for achieving context-sensitivity, which are very similar and almost interchangeable in formalizing the same concepts. The most obviously related techniques are ECA (event-condition-action) rules, production rules, decision trees, integrity constraints in abductive logic programs, and case-based reasoning. However, even the proposed uses of Description Logics [van Bunningen et al. 2006], Hierarchical Task Networks [Amigone et al. 2005], BDI-style commitment rules [Keegan et al. 2008]), and the agent cycles of da Silva and Vasconcelos [2007], for context-dependency and responsive environments, have much in common, and seem in fact, interchangeable. Other authors (e.g. Muniz et al. [2003]; Rodriguez et al. [2005]) use their own ad hoc formalizations, but these also bear remarkable similarities to ECA or production rules. Any significant differences among these techniques and their relative advantages and disadvantages may come to light in the future only when we consider richer requirements for formalization and reasoning, for example where temporal reasoning or default reasoning is crucial, where formal verification is attempted, or where there is a need for complex background theories to be used in conjunction with rules formalizing contexts.

Another fairly common feature of some of the systems reviewed in the survey is the recognition of the need for dynamic self-organization of devices within the AmI environment. Here again, a variety of different techniques is used to obtain similar functionality. These techniques are primarily based on architectures for communication among the agents that model the environment and the devices. Such architectures include, for example, the tuple space communication of da Silva and Vasconcelos [2007], the role-based communication of Busetta et al. [2003, 2004], the goal-based organization and interaction of Encarnacao and Kirste [2005], and the use of JINI in Amigone et al. [2005]. All these architectures have the primary aim of allowing agents to enter and leave the system, and for the goals of the system to be achieved by organizations of agents that form dynamically.

Learning is a prominent feature particularly in the smart home applications, with a variety of proposed techniques, for example reinforcement learning and data mining in MavHome [Cook et al. 2006] and fuzzy logic in iDorm [Hagras et al. 2004]. Learning also plays an implicit part in the recommender system applications, where, for example in Masthoff et al. [2007], the profile of the user in one domain is generalized and transferred to another domain.

There is much agreement about concerns over security and the social and ethical implications of AmI. There is clear agreement about the reasons why AmI gives rise to security concerns. The reasons include the collection of large amounts of personal data, the long-term persistence and integration of such data and the possibility of, and in fact often the need for, providing easy access to the data. Recent serious incidents of loss of data by institutions do not encourage optimism regarding security. In the UK alone, there have been recent losses and theft of personal data of large numbers of people (a figure of 3 million has been quoted) related to vehicle licensing, more than a hundred incidents, each involving loss or theft of a few thousand confidential patient health files, and the unlawful sale of thousands of personal records by a mobile phone company. To my knowledge no one has been held legally responsible for any of these losses of data, and there is no clear identification of who should be held responsible for the consequences suffered by individual victims of these losses or thefts. Undoubtedly, with AmI systems of the future these risks and problems will escalate.

We have all probably already had a taste of current infant AmI technologies in our everyday lives. Some of these experiences are undoubtedly very frustrating, such as sensor-operated taps that take much hand-waving before they produce a drop of water, sensor-operated lecture theatres that decide the blinds must be left open, no matter how many buttons one presses before a slide show, and heat-sensitive underfloor heating systems that have minds of their own. But arguably the potential benefits of AmI for individuals, institutions and businesses outweigh these initial frustrations and the security concerns. Moreover, the potential impact of AmI on and its challenges for research and development are undoubtedly immense and exciting.

ACKNOWLEDGMENTS

I am grateful to Bob Kowalski for very helpful comments and discussions. I am also grateful to the anonymous reviewers.

REFERENCES

AARTS, E. 2004. Ambient intelligence: a multimedia perspective. *IEEE Intell. Syst. 19*, 1, 12–19.

AARTS, E. H., AARTS, E. H. L., AND ENCARNAÇÃO, J. L. 2009. *True Visions: The Emergence of Ambient Intelligence*, Springer.

ABASCAL, J., BONAIL, B., MARCO, A., CASAS, R., AND SEVILLANO, J. L. 2008. AmbienNet: an intelligent environment to support people with disabilities and elderly people. In *Proceedings of the 10th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS)*. 293–294.

ABASCAL, J. 1997. Ethical and social issues of the teleservices, for disabled and elderly people. In *The Ethical Global Information Society. Culture and Democracy Revisited*, J. Berleur and D. Whitehouse, Eds. Chapman and Hall.

AGRAWAL, R., AND SRIKANT, R. 1995. Mining sequential patterns. In *Proceedings of the 11th International Conference on Data Engineering*, 3–14.

ALBRECHT, K. AND MCINTYRE, L. 2006. *Spychips: How Major Corporations and Government Plan to Track Your Every Move with RFID*. Plume (Penguin).

AMIGONE, F., GATTI, N., PINCIROLI, C., AND ROVERI, M. 2005. What planner for ambient intelligence applications? *IEEE Trans. Syst. Cybern. A Syst. Humans, 35* 1, 7–21.

AUGELLO, A., SANTANGELO, A., SORCE, S., PILATO, G., GENTILE, A., GENCO, A., AND GAGLIO, S. 2007. A multimodal interaction guide for pervasive services access. In *Proceedings of IEEE International Conference on Pervasive Services (PERSER)*, 250–256.

AUGUSTO, J. C., AND NUGENT, C. D. 2004. The use of temporal reasoning and management of complex events in Smart Homes. In *Proceedings of the European Conference on Artificial Intelligence (ECAI'04)*. R. L. de Mántaras and L. Saitta, Eds. IOS Press, 778–782.

AUGUSTO, J. C., NUGENT, C. D., AND BLACK, N. D. 2004. Management and analysis of time-related data in Smart Homes environments. In *Proceedings of the Mediterranean Conference on Medical and Biological Engineering (Medico)*, 4–7.

AUGUSTO, J. C., NUGENT, C. D., MARTIN, S., AND OLPHERT, C. 2005. Towards personalization of services and an integrated service model for Smart Homes applied to elderly care. In *Proceedings of the International Conference on Smart Homes and Health Telematics. Montreal, Canada, From Smart Homes to Smart Care*, S. Giroux and H Pigot, Eds. IOS Press. 151–158.

AUGUSTO, J. C., AND NUGENT, C. D. 2006a. Smart homes can be smarter. In *Designing Smart Homes: The Role of Artificial Intelligence*, J. C Augusto and C. D. Nugent, Eds., Lecture Notes in Artificial Intelligence, vol. 4008, Springer, Berlin, 1–15.

AUGUSTO, J. C., AND NUGENT, C. D. 2006b. *Designing Smart Homes The Role of Artificial Intelligence*, Lecture Notes in Artificial Intelligence, vol. 4008, Springer, Berlin.

AUGUSTO, J. C., LIU, J., MCCULLAGH, P., WANG, H., AND YANG, J.-B. 2008. Management of uncertainty and spatio-temporal aspects for monitoring and diagnosis in a smart home. *Inter. J. Computat. Intell. Syst., 1*, 4, 361–378.

BAADER, F., CALVANESE, D., MCGUINNESS, D., NARDI, D., AND PATEL-SCHNEIDER, P. 2003. *The Description Logic Handbook*. Cambridge University Press.

BAHADORI, S., CESTA, A., IOCCHI, L., LEONE, G. R., NARDI, D., PECORA, F., RASCONI, R., AND SCOZZAFAVA, L. 2004a. Towards ambient intelligence for the domestic care of the elderly, Robocare Tech. Rep. RC-TR-0704-3.

BAHADORI, S., CESTA, A., IOCCHI, L., LEONE, G. R., NARDI, D., PECORA, F., RASCONI, R., AND SCOZZAFAVA, L. 2004b. Towards ambient intelligence for the domestic care of the elderly. In *Ambient Intelligence: A Novel Paradigm*, P. Remagnino, G. L. Foresti, T. and Ellis, Eds., Springer.

BAHADORI, S., CESTA, A., GRISETTI, G., IOCCHI, L., LEONE, R., NARDI, D., ODDI, A., PECORA, F., AND RASCONI, R. 2004c. RoboCare: Pervasive intelligence for the domestic care of the elderly. *Intelligenza Artificiale, 1*, 1, 16–21.

BALLEGAARD, A. S., BUNDE-PEDERSEN, J., AND BARDRAM, J. E. 2006. Where to, Roberta?: reflecting on the role of technology in assisted living. In *Proceedings of the NordiCHI Conference*. 373–376.

BATSON, C., TURK, C., SHAW, L., AND KLEIN, T. 1995. Information function of emphatic emotion: learning that we value the other's welfare. *J. Personal. Social Psych. 68*, 300–313.

BELLIFEMINE, F., POGGI A., AND RIMASSA, G. 1999. JADE: A FIPA-compliant agent framework. In *Proceedings of The Practical Applications on Intelligent Agents and MultiAgent Technology (PAAM)*. 97–108, http://sharon.cselt.it/projects/jade/.

BJHC&IM. 2006. *The British Journal of Healthcare Computing & Information Management, 23*, 1, 39.

BLUM, A. AND FURST, M. 1995. Fast planning through graph analysis. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI'95)*. 1636–1642.

BOELLA, G., VAN DER TORRE, L., AND VILLATA, S. 2008. Institutional social networks for ambient intelligence, *AISB Convention: Communication, Interaction and Social Intelligence*.

BOHN, J., COROAMA, V., LANGHEINRICH, M., MATTERN, F., AND ROHS, M. 2004. Living in a world of smart everyday objects—Social, economic, and ethical implications. *Human Ecol. Risk Assess., 10*, 5.

BOSSE, T., JONKER, C. M., MEIJ, L., VAN DER SHARPANSKYTH, A., AND TREUR, J. 2006. Specification and verification of dynamics in cognitive agent models. In *Proceedings of the 6th International Conference on Intelligent Agent Technology (IAT'06)*, T. Nishida, et al. Eds., IEEE Computer Society Press, Los Alamitos, 247–254.

BOSSE, T., HOOGENDOORN, M., KLEIN, M. C. A., AND TREUR, J. 2008. A component-based ambient agent model for assessment of driving behaviour. In *Proceedings of the 5th International Conference on Ubiquitous Intelligence and Computing (UIC)*. F. E. Sandnes, Y. Zhang, C. Rong, L. T. Yang, J. Ma, Eds., Lecture Notes in Computer Science, vol. 5061, Springer, Berlin, 229–243.

Bouchard, B., Roy, P., Bouzouane, A., Giroux, S., and Mihailidis, A. 2008. An activity recognition model for Alzheimer's patients: extension of the COACH task guidance system. In *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI)*. IOS Press, Amsterdam, 2008, 811–812.

Brazier, F. M. T., Jonker, C. M., and Treur, J. 2000. Compositional design and reuse of a generic agent model. *Appl. Artif. Intell. J. 14*, 491–538.

Bresciani, P., Giorgini, P., Giunchiglia, F. Mylopoulos, J., and Prini, A. 2004. Tropos: an agent-orienyted software development methodology. *J. Auton. Agents Multiagent Syst., 8*, 3, 203–236.

Busetta, P., Merzi, M., Rossi, S., and Zanacanaro, M. 2003. Group communication for real-time role coordination and ambient intelligence. In *Proceedings of the Workshop on Artificial Intelligence in Mobile Systems (*in conjunction with *UbiComp'03)*.

Busetta, P., Kuflik, T., Merzi, M., and Rossi, S. 2004. Service delivery in smart environments by implicit organisations. In *Proceedings of the 1st International conference on Mobile and Ubiquitous Systems (MobiQuitous)*. IEEE Comp. Society.

Cai, Y. and Abascal, J. 2006. *Ambient Intelligence for Everyday Life*. Vol. 3864. Springer, Berlin.

CASPIAN (Consumers Against Supermarket Privacy and Numbering). http://www.nocards.org/.

Cesta, A., Oddi, A., and Smith, S. F. 2002. A constrained-based method for project scheduling with time windows. *J. Heuristics, 8*, 1, 109–135.

Cesta, A., Cortellessa, G., Pecora, F., and Rasconi, R. 2005. Exploiting scheduling techniques to monitor the execution of domestic activities. *Intelligenza Artificiale, 2*, 4.

Chervest, K., Davies, N., and Mitchell, K. 2002. The role of adaptive hypermedia within a context-aware tourist guide. *Comm. ACM, 45*, 5, 47–51.

Collier, R. W., O'Hare, G. M. P., Lowen, T., and Rooney, C. F. B. 2003. Beyond prototyping in the factory of agents. In *Proceedings of the 3rd International Central & Eastern European Conference on Multi-Agent Systems (CEEMAS)*. Lecture Notes in Computer Science, Springer, Berlin.

Cook, D. J., Youngblood, M., and Das, S. K. 2006. A multi-agent approach to controlling a smart environment. In *Designing Smart, The Role of Artificial Intelligence*, J. C Augusto and C.D. Nugent, Eds., Lecture Notes in Artificial Intelligence, vol. 4008, Springer, Berlin, 165–206.

Corchado, J. M., Bajo, J., and Abraham, A. 2008. GerAmi: improving healthcare delivery in geriatric residences. *J. IEEE Intell. Syst.* (Special Issue on Ambient Intelligence), *3*, 2, 19–25.

Coroama, V. and Rothenbacher, F. 2003. The chatty environment—providing everyday independence to the visually impaired. In *Proceedings of the 2nd International Workshop on Ubiquitous Computing for Pervasive Healthcare Applications (UbiHealth)*.

Costantini, S., Mostarda, L., Tocchio, A., and Tsintza, P. 2008. DALICA: Agent-based ambient intelligence for cultural-heritage scenarios. *IEEE Intell. Syst.* (Special Issue on Ambient Intelligence), *23*, 2, 34–41.

Cowie, R., Douglas-Cowie, E., Tsapatsoulis, N., Votsis, G., Kollias, S., Fellenz, W., and Taylor, J. 2001. Emotion recognition in human-computer interaction. *IEEE Signal Process. Mag. 18*, 1, 32–80.

Craig, S. D., Graesser, A. C., Sullins, J., and Gholson, B. 2004. Affect and learning: an exploratory look into the role of affect in learning. *J. Educat. Media, 29*, 241–250.

da Silva, F. S. C. and Vasconcelos, W. W. 2007. Managing responsive environments with software agents. *J. Appl. Artif. Intell. 21*, 4, 469–488.

da Silva, F. S. C. and Vasconcelos, W. W. 2005. Agent-Based Management of Responsive Environments. In *Proceedings of Advances in Artificial Intelligence, 9th Congress of the Italian Association for Artificial Intelligence*, 224–236.

Damasio, A. R. 2000. *The Feeling of What Happens-Body, Emotion and the Making of Consciousness*, Harvest Books, London, UK.

Desjardins, M., Durfee, E. H., and Wolverton, M. 1999. A survey of research in distributed, continual planning. *AI Mag. 20*, 4, 13–22.

D'Mello, S., Jackson, T., Craig, S., Morgan, B., Chipman, P., White, H., Person, N., Kort, B., el Kaliouby, R., Picard., R. W., and Graesser, A. 2008. AutoTutor detects and responds to learners affective and cognitive states. In *Proceedings of the Workshop on Emotional and Cognitive Issues at the International Conference of Intelligent Tutoring Systems*.

Ducatel, K., Bogdanowicz, M., Scapolo, F., Leijten, J., and Burgelman, J.-C. 2001. Scenarios for ambient intelligence in 2010. IST Advisory Group Final Report, European Commission, EC. Brussels.

Ducatel, K., Bogdanowicz, M., Scapolo, F., Leijten, J., and Burgelman, J.-C. 2003. Ambient intelligence : From vision to reality. IST Advisory Group Draft Report, European Commission.

Ekman, P. 1992. An argument for basic emotions. *Cognition Emotion, 6*, 4.

ENCARNACAO, J. L. AND KIRSTE, T. 2005. Ambient intelligence: Towards smart appliance ensembles. M. Hemmje et al. Eds., Lecture Notes in Computer Science, vol. 3379, Sprinter, Berlin, 261–270.

ENDRISS, U., MANCARELLA, P., SADRI, F., TERRENI, G., AND TONI, F. 2004. The CIFF proof procedure for abductive logic programming with constraints. In *Proceedings of the International Conference on Logics in AI (JELIA'04)*. Lecture Notes in Artificial Intelligence, Springer, Berlin, 31–43.

FAVELA, J., RODRIGUEZ, M., PRECIADO, A., AND GONZALEZ, V. M. 2004. Integrating context-aware public displays into a mobile hospital information system. *IEEE Trans. Inf. Technol. Biomed. 8*, 3, 279–286.

FENG, L., APERS, P. M. G., AND JONKER, W. 2004. Towards context-aware data management for ambient intelligence, In *Proceedings of the 15th International Conference on Database and Expert Systems Applications (DEXA)*. Lecture Notes in Computer Science, vol. 3180, 422–431.

FIFER, S. K. AND THOMAS, N. 2002. A second and more promising round of e-health enterprises. *Manag. Care Interface, 15*, 52–56.

FIRBY, R. J. 1999. *The RAPS Language Manual*. Neodesic, Inc. Chicago, IL.

FOREST, F., OEHME, A., YAICI, K., AND VERCHERE-MORICE, C. 2006. Psycho-social aspects of context awareness in ambient intelligent mobile systems. In *Proceedings of the IST Summit Workshop, Capturing Context and Context Aware Systems and Platforms*.

FRIEDEWALD, M., DA COSTA, O., PUNIE, Y., ALAHUHTA, P., AND HEINONEN, S. 2005. Perspectives of ambient intelligence in the home environment. *Telematics Informatics, 22*, Elsevier, 221–238.

GAGGIOLI, A. 2005. Optimal experience in ambient intelligence. In *Ambient Intelligence*, G. Riva, F. Vatalaro, F. Davide, and M. Alcaniz, Eds., IOS Press, Amsterdam, 35–43.

GALTON, A. 2004. Eventualities. In *The Handbook of Time and Temporal Reasoning in Artificial Intelligence*, Vila, Van Beek, Boddy, Fisher, Gabbay, Galton, and Morris, Eds., MIT Press.

GARATE, A., HERRASTI, N., AND LOPEZ, A. 2005. GENIO: An ambient intelligence application in home automation and entertainment environment. In *Proceedings of the Joint sOc-EUSAI Conference*. 241–245.

GEIB, C. W. AND GOLDMAN, R. P. 2001. Probabilistic plan recognition for hostile agents. In *Proceedings of the 14th International Florida Artificial Intelligence Research Society Conference*. AAAI Press, 580–584.

GEIB, C. W. 2002. Problems with intent recognition elder care. In *Proceedings of the AAAI Workshop on Automation as Caregiver*.

GEORGEFF, M., LANSKY, A., AND BESSIER, P. 1985. A procedural logic. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence (IJCAI).*

GOMEZ, R., AUGUSTO, J. C., AND GALTON, A. 2001. Testing an event specification language. In *Proceedings of the 13th International Conference of Software Engineering and Knowledge Engineering (SEKE)*. 341–346.

GONZÁLEZ, G., LÓPEZ, B., AND DE LA ROSA, J. L. L. 2004a. Managing emotions in smart user models for recommender systems. In *Proceedings of 6th International Conference on Enterprise Information Systems (ICEIS). 5*, 187–194.

GONZÁLEZ, G., LÓPEZ, B., AND DE LA ROSA, J. L. L. 2004b. Smart user models for tourism: A holistic approach for personalised tourism services. *Inform. Technol. Tourism J.: Applicat. Methodol. Techniques, 6*, 4, Cognizant Communication Corporation.

GONZÁLEZ, G., CECILIO, A., LOPEZ, B., AND DE LA ROSA, J. L. L. 2005. Smart user models for ambient recommender systems. In *Ambient Intelligence and (Everyday) Life*, Y. Cai, Y. and J. Abascal, Eds., University of Basque Country, San Sebastian, Spain, 113–122.

GONZÁLEZ, G., DE LA ROSA, J. L. L., DUGDALE, J., PAVARD, B., EL JED, M., PALLAMIN, N., ANGULO, C., AND KLANN, M. 2006. Towards ambient recommender systems: results of new cross-disciplinary trends. In *Proceedings of ECAI Workshop on Recommender Systems*.

GORDON, T. F. AND KARACAPILIDIS, N. 1997. The Zeno argumentation framework, In *Proceedings of the 6th International Conference on Artificial Intelligence and Law*, ACM Press, New York, 10–18.

GOUAUX, F., SIMON-CHAUTEMPS, L., FAYN, J., ADAMI, S., ARZI, M., ASSANELLI, D., FORLINI, M.C., MALOSSI, C., MARTINEZ, A., PLACIDE, J., ZILIANI, G.L., AND RUBEL, P. 2002. Ambient intelligence and pervasive systems for monitoring of citizens at cardiac risk: New solutions from the EPI-MEDICS project. *Comput. Cardiol. 29*, 289–292.

GRUDIN, J. 2002. Group dynamics and ubiquitous computing. *Comm. ACM, 45*, 12.

GURALNIK, V., AND HAIGH, K. Z. 2002. Learning models of human behaviour with sequential patterns. In Proceedings of the *AAAI Workshop on Automation as Caregiver*. AAAI Press.

GUREVYCH, I. AND MÜHLHÄUSER, M. 2007. Natural language processing for ambient intelligence. *Special Issue of KI-Zeitschrift Ambient Intelligence und Künstliche Intelligenz*, 10–16.

GUREVYCH I. AND MÜHLHÄUSER, M. Eds. 2009. *Int. J. Seman. Comput.* (Special Issue on Ambient Computing).

HAGRAS, H., CALLAGHAN, V., COLLEY, M., CLARKE, G., POUNDS-CORNISH, A., AND DUMAN, H. 2004. Creating an ambient-intelligence environment using embedded agents. *IEEE Intell. Syst.* 12–20.

HAIGH, K. Z., AND KIFF, L. M. 2004. The independent lifestyle assistant (I.L.S.A.): AI lessons learned. In *Proceedings of the 6th Conference on Innovative Applications of Artificial Intelligence (IAAI)*.

HAIGH, K. Z., KIFF, L. M., AND HO, G. 2006. The Independent LifeStyle Assistant (I.L.S.A.): Lessons learned. *Assistive Technol. 18*, 1, 87–106.

HAUX, R. 2006. Individualization, globalisation and health—about sustainable information technologies and the aim of medical informatics. *Int. J. Medical Informatics 75*, Elsevier, 795–808.

HEINONEN, S. 2004. Mobile telework at the crossroads of social, environmental and cultural challenges. In *Proceedings of the 9th International Telework Workshop*.

HEINSOHN, J. 1994. Probabilistic description logics. In *Proceedings of Uncertainty in Artificial Intelligence*. 311–318.

HERBON, A., OEHME, A., AND ZENTSCH, E. 2006. Emotions in ambient intelligence—an experiment on how to measure affective states. In *Proceedings of the 20th British Computer Society HCI Group Conference*.

JIANG, X., HONG, J. I., AND LANDAY, J. A. 2002. Approximate information flows: socially-based modelling of privacy in ubiquitous computing. In *Procedings of the 4th International Conference on Ubiquitous Computing (UbiComp)*, Springer-Verlag, 176–193.

KAKAS, A. C., KOWALSKI, R., AND TONI, F. 1992. Abductive logic programming . *J. Logic Computation, 2*, 6, 719–770.

KARACAPILIDIS, N. AND PAPADIAS, D. 2001. Computer supported argumentation and collaborative decision making: The Hermes system. *Inform. Syst. 26*, 4, 259–277.

KEEGAN, S., O'HARE, G. M. P., AND O'GRADY, M. J. 2008. Easishop: ambient intelligence assists everyday shopping. *J. Inform. Sci. 178*, 3, 588–611.

KEMPPAINEN, E., ABASCAL, J., ALLEN, B., DELAITRE, S., GIOVANNINI, C., AND SOEDE, M. 2007. Ethical and legislative issues with regard to Ambient Intelligence. In *Towards an Inclusive Future. Impact and Wider potential of Information and Communication Technologies*, P. R. Roe, Ed. 188–205. http://www.tiresias.org/cost219ter/inclusive_future/inclusive_future_book.pdf

KIM, K. H., BANG, S. W., AND KIM, S. R. 2004. Emotion recognition system using short-term monitoring of physiological signals. *Medical Biological Engin. Comput. 42*, 419–427.

KOFOD-PETERSON, A., AND AAMODT, A. 2006. Contexualised ambient intelligence through case-based reasoning. In *Proceedings of the 8th European Conference on Case-Based Reasoning (ECCBR)*. T. R. Roth-Berghofer et al. Eds., Lecture Notes in Artificial Intelligence, vol. 4106, Springer-Verlag, Berlin, 211–225.

KOPACSI, S., KOVACS, G., ANUFRIEV, A., AND MICHELINI, R. 2007. Ambient intelligence as enabling technology for modern business paradigms. *Robotics Comput.-Integrat. Manufactur. 23*, 247–256.

KORHONEN, I., PARKKA, J., AND GILS, M. V. 2003. Health monitoring in the home of the future. *IEEE Eng. Med. Biol. Mag. 22*, 3, 66–73.

KWON, O. B., AND SADEH, N. 2004. Applying case-based reasoning and multi-agent systems to context-aware comparative shopping, *Decision Support Syst. 37*, 199–213.

LEVINSON, R. 1997. The planning and execution assistant and trainer (PEAT). *J. Head Trauma Rehab. 12*, 85–91.

LI, Y., FENG, L., AND ZHOU, L. 2008. Context-aware database querying: Recent progress and challenges. In *Context-Aware Mobile and Ubiquitous Computing for Enhanced Usability: Adaptive Technologies and Applications,* D. Stojanovic, Ed., IGI, 147–170.

LIU, J., AUGUSTO, J. C., AND WANG, H. 2006. Considerations on uncertain spatio-temporal reasoning in smart home systems. In *Proceedings of the 7th International Conference on Applied Artificial Intelligence*.

MA, T., KIM, Y. D., MA, Q., TANG, M., AND ZHOU, W. 2005. Context-aware implementation based on CBR for smart home. In *Proceedings of the IEEE International Conference on Wireless and Mobile Computing, Networking and Communications 2005 (WiMob)*. IEEE Computer Society, Vol. 4, 112–115.

MAHILIDIS, A., FERNIE, G., AND BARBENEL, J. C. 2001. The use of artificial intelligence in the design of an intelligent cognitive orthotic for people with dementia. *Assistive Technol. 13*, 23–39.

MAHILIDIS, A., CARMICHAEL, B., BOGER, J., AND FERNIE, G. 2003. An intelligent environment to support aging-in-place, safety, and independence of older adults with dementia. In *Proceedings of the 2nd International Workshop on Ubiquitous Computing for Pervasive Healthcare Applications (UbiHealth'03)*. J. McCarthy and J. Scott, Eds., 294–297.

MAHILIDIS, A., CARMICHAEL, B., AND BOGER, J. 2004. The use of computer vision in an intelligent environment to support aging-in-place, safety, and independence at home. *IEEE Trans. Inform. Technol. Biomed. 8*, 238–247.

MANCARELLA, P., SADRI, F., TERRENI, G., AND TONI, F. 2004. Planning partially for situated sgents. In *Proceedings of 5th Workshop on Computational Logic in Multi-Agent Systems (CLIMA V)*. J. Leite and P.Torroni, Eds., Springer Verlag, 29–30.

MANCARELLA, P., TERRENI, G., SADRI, F., TONI, F., AND ENDRISS, U. 2009. The CIFF proof procedure for abductive logic programming with constraints: Theory, implementation and experiments. *Theory Pract. Logic Program. 9*, 6, 691–750.

MARCHESOTTI, L., PIVA, S., AND REGAZZONI, C. 2005. Structured context-analysis techniques in biologically inspired ambient-intelligence systems. *IEEE Trans. Syst. Man, Cybernet. A, Syst. Humans. 3*, 1, 106–120.

MARREIROS, G., SANTOS, R., RAMOS, C., NEVES, J., NOVAIS, P., MACHADO, J., AND BULAS-CRUZ, J. 2007a. Ambient intelligence in emotion based ubiquitous decision making. In *Proceedings of the 2nd Workshop on Artificial Intelligence Techniques for Ambient Intelligence (AITAmI)*, J. C. Augusto and D. Shapiro, Eds., 86–91.

MARREIROS, G., SANTOS, R., NOVAIS, P., MACHADO, J., RAMOS, C., NEVES, J., AND BULAS-CRUZ, J. 2007b. Argumentation-based decision making in ambient intelligence environments. In *Proceedings of Progress in Artificial Intelligence, 13th Portuguese Conference on Artificial Intelligence (EPIA)*, Lecture Notes in Computer Science, Vol. 4874, Springer Berlin, 309–322.

MASTHOFF, J., AND GATT, A. 2006. In pursuit of satisfaction and prevention of embarrassment: affective state in group recommender systems. *User Model. User-Adapt. Interac. J. Personal. Res. 16*, 281–319.

MASTHOFF, J., VASCONCELOS, W. W., AITKEN, C., AND CORREA DA SILVA, F. S. 2007. Agent-based group modelling for ambient intelligence. In *Proceedings of the AISB Symposium on Affective Smart Environments*. P. Oliver and C. Kray, Eds. 90–96.

MCLAUGHLIN, A. C., ROGERS, W. A., AND FISK, A. D. 2002. Effectiveness of audio and visual training presentation modes for glucometer calibration, In *Proceedings of the Human Factors and Ergonomics Society 46th Annual Meeting*, 2059–2063.

MISKER, J. M. V., VEENMAN, C. J., AND ROTHKRANTZ, L. J. M. 2004. Groups of collaborating users and agents in ambient intelligent environments. In *Proceedings of the 3rd International Conference on Autonomous Agents and Multi Agent Systems (AAMAS)*. C. Sierra and L. Sonenberg, Eds., IEEE Computer Society, 1320–1321.

MITCHELL, T. 1997. Decision tree learning. In *Machine learning*, T. Mitchell, Ed., The McGraw-Hill Companies, Inc., 52–78.

MOZER, M. C. 2005. Lessons from an adaptive house. In *Smart Environments: Technologies, Protocols, and Applications*, D. Cook and R. Das, Eds., Wiley & Sons, Hoboken, New Jersey, 273–294.

MUNOZ, M. A., RODROGUEZ, M., FAVELA, J., MARTINEZ-GARCIA, A. I., AND GONZALEZ, V. M. 2003. Context aware mobile communication in hospitals. *IEEE Comput. 36*, 38–46.

MYKITYSHYN, A. L., FISK, A. D., AND ROGERS, W. A. 2002. Learning to use a home medical device: mediating age-related differences with training. *Human Factors, 44*, 354–364.

NIEMELA, M., FUENTETAJA, R. G., KAASINEN, E., AND GALLARDO, J. L. 2007. Supporting independent living of the elderly with mobile-centric ambient intelligence: user evaluation of three scenarios, B. Schiele et al. Eds., Lecture Notes in Computer Science, vol. 4794, Springer-Verlag Berlin, 91–107.

NORWOOD, M. 2000. Affective feedback devices for continuous usability assessment. Masters thesis, Media Arts and Sciences, MIT, Cambridge, MA.

OATLEY, K., AND JOHNSON-LAIRD, P. N. 1987. Towards a cognitive theory of emotions. *Cognition Emotion, 1*, 29–50.

PAIVA, A., PRADA, R., AND PICARD, R. W. 2007. Affective computing and intelligent interaction. In *Proceedings of the 2nd International Conference on Affective Computing and Intelligent Interaction, (ACII'07)*. 2007, A. Paiva, R. Prada, and R.W. Picard, Eds., Lecture Notes in Computer Science, vol. 4738, Springer Verlag, Berlin.

PARK, D., NAM, T., SHI, C., GOLUB, G., AND VAN LOAN, C. 2006. Designing an immersive tour experience system for cultural tour sites. In *Proceedings of Extended Abstracts CHI.* ACM Press, New York, 1193–1198.

PATKOS T., BIKAKIS, A., ANTONIOU, G., PAPADOPOULI, M., AND PLEXOUSAKIS, D. 2007. Distributed AI for ambient intelligence: issues and approaches, In *Proceedings of the European Conference on Ambient Intelligence (AmI)*. B. Schiele, A. K. Dey, H. Gellersen, B. de Ruyter, M. Tscheligi, R. Wichert, E. Aarts, and A. Buchmann (Eds.), Lecture Notes in Computer Science, Vol. 4794, Springer Berlin, 156–176.

PATON, N., AND DIAZ, O. 1999. Active database systems. *ACM Comput. Surv. 31*, 1, 63–103.

PENSERINI, L., BRESCIANI, P., KUFLIK, T., AND BUSETTA, P. 2005. Using Tropos to model agent based architectures for adaptive systems: A case study in ambient intelligence. In *Proceedings of the IEEE International Conference on Software—Science, Technology and Engineering (SwSTE)*. IEEE Press.

PETERSEN, S. A. AND KOFOD-PETERSEN, A. 2006. The non-accidental tourist: Using ambient intelligence for enhancing tourist experiences. In *Network-Centric Collaboration and Supporting Frameworks*, L. Camarinha-Matos, H. Afsarmanesh, and M. Ollus, Eds., International Federation for Information Processing (IFIP). Vol. 224, 619–626.

PHILIPOSE, M., FISHKIN, K. P., PERKOWITZ, M., PATTERSON, D. J., FOX, D., KAUTZ, H., AND HAHNEL, D. 2004. Inferring activities from interactions with objects. *Pervasive Comput. 3*, 50–57.

PICARD, R. W. 2000. Towards computers that recognise and respond to user emotion, Tech. Rep. 0018-8670/00 2000 IBM.

PICARD, R. W. 2007. Toward machines with emotional intelligence. In *The Science of Emotional Intelligence: Knowns and Unknowns*, G. Matthews, M. Zeidner, and R.D. Roberts, Eds., Oxford University Press, Oxford, UK.

PIVA, S., BONAMICO, C., REGAZZONI, C., LAVAGETTO, F. 2005. A flexible architecture for ambient intelligence systems supporting adaptive multimodal interaction with users. In *Ambient Intelligence*, G. Riva, F. Vatalaro, F. Davide, and M. Alcañiz, Eds., IOS Press, http://www.ambientintelligence.org.

PLOCHER, T. AND KIFF, L. M. 2003. Mobility monitoring with the Independent LifeStyle Assistant (I.L.S.A.) . In *Proceedings of the International Conference on Aging, Disability and Independence (ICADI)*. 170–171.

POLLOCK, M. E, BROWN, L., COLBRY, D., MCCARTHY, C. E., OROSZ, C., PEINTNER, B., RAMAKRISHNAN, S., AND TSAMARDINOS, I. 2003. Autoreminder: an intelligent cognitive orthotic system for people with memory impairment. *Robotics Auton. Syst. 44*, 273–282.

PRAKKEN, H. AND GORDON, T. F. 1999. Rules of order for electronic group decision making - a formalization methodology. In *Proceedings of the VIM Spring and Winter Workshops on Collaboration between Human and Artificial Societies*. Lecture Notes in Artificial Intelligence, vol. 1624, Springer Verlag, Berlin, 246–263.

RAMOS, C., AUGUSTO, J. C., AND SHAPIRO, D. Eds. 2008. *IEEE Intel. Syst. Mag. 23*, 2,

RAMOS, C. 2008. An architecture for ambient intelligence environments. In *Advances in Soft Computing, 51*, Springer, 30–38.

RAO, A. S. AND GEORGEFF, M. P. 1991. Modelling rational agents within a BDI architecture. In *Principles of Knowledge Representation and Reasoning*.

RAO, A. S. AND GEORGEFF, M. P. 1995. BDI-agents: From theory to practice. In *Proceedings of the 1st International Conference on Multiagent Systems (ICMAS)*. 312–319.

RIVA, G. 2003. Ambient intelligence in health care. *CyberPsych. Behav. 6*, 3, Mary Ann Liebert, Inc. Publishers, 295–300.

RODRIGUEZ, M. D., FAVELA, J., MARTINEZ, E. A., AND MUNOZ, M. A. 2004. Location-aware access to hospital information and services. *IEEE Trans. Inform. Technol. Biomed. 8*, 4, 448–455.

RODRIGUEZ, M., FAVELA, J., PRECIADO, A., AND VIZCAINO, A. 2005. Agent-based ambient intelligence for healthcare. *AI Comm. 18*, 3, 201–216.

ROUVROY, A. 2008. Privacy, data protection, and the unprecedented challenges of ambient intelligence. *Studies Ethics, Law, Technol. 2*, 1, Article 3.

RUMETSHOFER, H., PUHRETMAIR, F., AND WOB, W. 2003. Individual information presentation based on cognitive styles for tourism information systems, In *Proceedings of the International Conference on Information and Communication Technologies in Tourism*. A. J. Frew, M. Hitz, and P. O'Connor, Eds., Springer Verlag, Berlin, 440–449.

SACERDOTI, E. 1977. *A Structure for Plans and Behaviour*. North Holland, Elsevier.

SADRI, F. AND TONI, F. 2000. Abduction with negation as failure for active and reactive rules. In *Proceedings of the 6^{th} Congress of the Italian Association for Artificial Intelligence*. Lecture Notes in Artificial Intelligence, vol. 1792, E. Lamma and P. Mello, Eds., 49–60.

SADRI, F., TONI, F., AND TORRONI, P. 2002. Resource reallocation via negotiation through abductive logic programming. In *Proceedings of the International Conference on Logics in AI*. Lecture Notes in Artificial Intelligence, vol. 2424, Springer-Verlag, 419–431.

SADRI, F. 2007. Ambient Intelligence for Care of the Elderly in Their Homes. In *Proceedings of the 2nd Workshop on Artificial Techniques for Ambient Intelligence (AITAmI)*. 62–67.

SADRI, F. AND STATHIS, K. 2008. Ambient intelligence. In *Encyclopedia of Artificial Intelligence*, J.R. Rabunal, J. Dorado, and A. P. Sierra, Eds., IGI Global.

SAFFIOTTI, A. AND BROXVALL, M. 2005. PEIS ecologies: ambient intelligence meets autonomous robotics. In *Proceedings of the Joint Conference on Smart Objects and Ambient Intelligence: Innovative Context-Aware Services: Usages and Technologies*. ACM International Conference Proceeding Series, volume 121, 277–281.

SALOVEY, P. AND MAYER, J. D. 1990. Emotional intelligence. *Imagination, Cognition, Personality 9*, 3, 185–211.

SÁNCHEZ, D., TENTORI, M., AND FAVELA, J. 2008. Activity recognition for the smart hospital, *J. IEEE Intell. Syst.* (Special Issue on Ambient Intelligence) *23*, 2, 50–57.

SCHANK, R. 1982. *Dynamic Memory: A Theory of Learning in Computers and People*. Cambridge University Press, Cambridge, UK.

SEEM. 2003. Single European market, Rep. of SEEM Workshop.

SIMPSON, R., SCHRECKENGHOST, D., LOPRESTI, E. F., AND KIRSCH, N. 2006. Plans and planning in smart homes. In *Designing Smart Homes The Role of Artificial Intelligence*, J. C Augusto and C. D. Nugent, Eds., Lecture Notes in Artificial Intelligence, vol. 4008, Springer Verlag, Berlin, 71–84.

SLEETH, C. E. 2002. Technology map: Pervasive computing. Tech. rep. SRI Consulting Business Intelligence.

STANKOVSKI, V. AND TRNKOCZY, J. 2006. Application of decision trees to smart homes. In *Designing Smart Homes The Role of Artificial Intelligence*, J. C. Augusto and C. D. Nugent, Eds., Lecture Notes in Artificial Intelligence, vol. 4008, Springer Verlag, Berlin, 132–145.

TATE, A, DALTON, J., AND LEVINE, J. 2000. O-Plan: a web-based AI palnning agent. In *Proceedings of the 7th National Conference on Artificial Intelligence (AAAI)*. AAAI Press, 1131–1132.

THACKARA, J. 2001. The design challenge of pervasive computing. *Interactions, 8*, 3, 46–52.

UN 2002. Population aging. UN Department of Economic Social Affairs, Population Division. http://www.un.org/esa/population/publications/ageing/Graph.pdf .

UN 2004. World population to 2300. UN Department of Economic Social Affairs, Population Division, United Nations, New York. http://www.un.org/esa/population/publications/longrange2/WorldPop2300final.pdf.

VAN BUNNINGEN, A. H., FENG, L., AND APERS, P. M. G. 2006. A context-aware query preference model for ambient intelligence. *In Proceedings of the 17th International Conference on Database and Expert Systems Applications (DEXA)*. 33–43.

VAN BUNNINGEN, A. H., FOKKINGA, M., APERS, P., AND FENG, L. 2007. Ranking query results using context-aware preferences. In *Proceedings of the 1st International Workshop on Ranking in Databases (In Conjunction with ICDE)*. IEEE Computer Society Press. http://www.cs.uwaterloo.ca/conferences/dbrank/2007/paper21/paper.pdf.

VAN HEERDE, H., ANCIAUX, N., FENG, L., AND APERS, P. M. G. 2006. Balancing smartness and privacy for ambient intelligence. In *Proceedings of the 1st European conference on Smart Sensing and Context (EuroSSC)*, 255–258.

VASTFJALL, D., FRIMAN, M., GARLING, T., AND KLEINER, M. 2000. The measurement of core affects: a Swedish self-report measure derived from the affect circumplex. *Goteborg Psychol. Rep. 30*, 8, Goteborg University.

WEBER, W. 2003. Ambient intelligence - Industrial research on a visionary concept, In *Proceedings of the International Symposium on Low Power Electronics and Design (ISLPED)*. 247–251.

WERNER, W., RABAEY, J. M., AND AARTS, E. H. L. 2005. *Ambient Intelligence*. Springer.

WEISER, M. 1991. The computer for the 21st century. In *Scientific Amer.* 94-104.

WEISER, M. 1993. Some computer science issues in ubiquitous computing. *Comm. ACM 36*, 7, 75–85.

WORLD HEALTH ORGANISATION. 2005. eHealth resolution. In *Proceedings of the 8th World Health Assembly.* WHO, Geneva. http://www.who.int/gb/e/e_wha58.html.

WRIGHT, D., GUTWIRTH, S., FRIEDEWALD, M., VILDJIOUNAITE, E., AND PUNIE, Y. Eds. 2008. *Safeguards in a World of Ambient Intelligence*, The International Library of Ethics, Law and Technology, Springer.

YANG, J., LIU, J., WANG, J., SII, H., AND WANG, H. 2006. Belief rule-based inference methodology using the evidential reasoning approach – RIMER, IEEE *Trans. Syst. Man, Cybernetics, Part A, 36*, 2, 266–285.

ZIV, J. AND LAMPEL, A. 1978. Compression of individual sequences via variable rate coding. *IEEE Trans. Inform. Theory, 24*, 530–536.

Received April 2009; revised February 2010; accepted February 2010