Presented at The 8th International Conference on Intelligent Environments - IE'12, Guanajuato, Mexico, 26-29 June 2012

# Explorations of Autonomy

## An Investigation of Adjustable Autonomy in Intelligent Environments

Matthew Ball[1], Vic Callaghan
Intelligent Environments Group
School of Computer Science and Electronic Engineering
University of Essex, Colchester, UK.
[1]mhball@essex.ac.uk

***Abstract*—There are many arguments for and against the use of autonomous-agents in intelligent environments. Some researchers maintain that it is of utmost importance to give complete control to users, and hence greatly restrict autonomy of agents; whereas, others believe that is it preferable to increase user convenience by allowing agents to operate autonomously on the user's behalf. While both of these approaches have their distinct merits, they are not suitable for all users. As people's opinions and concerns regarding agent autonomy are highly individual, depending on a wide range of factors and often changing over time, a much more dynamic approach to agent autonomy is needed. This work explores how it is possible to equip intelligent environments with an adjustable autonomy mechanism, which allows an individual user to increase or decrease agent autonomy in order to find their own comfortable sweet-spot between maintaining/relinquishing control and gaining/losing convenience. This paper presents the Adjustable Autonomy Intelligent Environment (AAIE) model, discusses how adjustable autonomy can be achieved in intelligent environments, and discusses the major findings from a recent online survey and user study, which highlight the major factors and concerns of users that determine their personal preferences towards different levels of autonomy.**

***Keywords-intelligent environments; ambient intelligence; autonomous-agents; adjustable autonomy; pervasive computing;***

## I. INTRODUCTION

For many years people have designed and produced systems with the aim of automation – creating systems to perform tasks on behalf of people when the task at hand is perhaps too tedious, dangerous or difficult. As technology has developed, the area of artificial intelligence grew and the paradigm of agents was created. With these agents we are able to create technology that not only acts on our behalf but acts with a known purpose or goal, and can reason about its actions, learn new actions, and even in some instances refuse to take action. This has allowed a leap forward from automatic systems to autonomous-agent systems – not only acting by themselves but also governing their actions and adapting how they operate [1].

The use of autonomous-agents in intelligent environments has been a much debated topic. Some believe that there is a risk of creating something comparable to Bentham's Panopticon [2] or some notion of 'Big Brother' being able to monitor our every move and know all of our personal interests, as in the famous book Nineteen Eighty-Four by George Orwell [3]. Moreover, as we're told by Callaghan et al., an end-user driven approach to intelligent environment management can encourage creativity in users, it goes beyond the "current DIY approach of paint and wallpaper" and allows people to customise (or decorate) their homes in a digital sense [4]. For reasons such as these, many researchers in intelligent environments take the stance that the use of autonomous-agents should be greatly restricted, and instead the end-user should always be given complete control over all systems. In these end-user driven approaches, it becomes the responsibility of the user to program the intelligent environment in order to create automated behaviours, although the user of the system may not actually have any knowledge of computer programming nor any technical knowledge of the system. An end-user driven approach usually adopts a simplified programming interface to enable the end-user to program the behaviour rules more easily, as in [5-7]. In most situations, producing a system that empowers the user might seem the logical choice; however, problems can arise in an end-user driven system since the intelligence and adaptability of the system depends heavily on the creativity, intelligence, willingness and ability of the user. For example, users may be too busy at times or may have a low level of confidence in their ability to manage such a complex system, or users may even have a physical disability and find it very difficult or even impossible to interact with computer devices. In these situations, autonomous-agents can be very useful as they are designed to operate on the user's behalf and greatly reduce the cognitive load, and sometimes the physical requirements, placed on the user in managing the intelligent environment [8]. In high level terms, autonomous-agent driven intelligent environments can be defined as those that employ artificial intelligence and machine learning mechanisms to program automated behaviour in the environment by monitoring and learning from the user's behaviours and interactions with the environment and system, as in [9, 10].

While both the end-user driven and autonomous-agent driven approaches have great advantages, they are only suited to certain types of users [11-13]. This work explores how we can make intelligent environments more dynamic and personalisable by equipping them with adjustable autonomy, and allowing the user to explore the trade-off between the convenience offered by autonomous-agents and the amount of control offered by end-user driven systems.

An online survey was recently conducted, which aimed to investigate people’s opinions of the use of autonomy in intelligent environments [14]. As a follow up to the online survey, a working adjustable autonomy intelligent environment has been implemented and a series of user trials were conducted, which aimed to gain deeper insights into the reasoning behind people’s attitudes of different levels of autonomy and explore how using adjustable autonomy can change people’s opinions of intelligent environments [14]. The major findings from these studies highlight the main factors and concerns of users that determine their attitudes and preferences towards agent autonomy in intelligent environments.

The remainder of this paper is structured as follows: Section II discusses the concept of autonomy. Section III describes how adjustable autonomy can be applied in intelligent environments and presents the Adjustable Autonomy Intelligent Environment (AAIE) model. Section IV discusses the major determining factors of users' autonomy preferences, and Section V gives a concluding discussion. This work aims to raise awareness of the issues related to using static (and extreme) levels of autonomy amongst researchers of intelligent environments and ambient intelligence systems.

## II. What is Autonomy?

In our lives as we grow and learn we become more autonomous, losing the dependences we have on others. Many people will remember the day they passed their driving test and got their first car, and the sense of freedom and independence that followed. This is a prime example of when someone gains autonomy in their lives. Up until this point the person was dependant on getting lifts from others or relying on public transport to get them were they need to go. With their own car, however, they are much more self-reliant (autonomous) and are able to travel wherever and whenever they decide.

The notion of autonomy is strongly related to independence. The Oxford English Dictionary defines autonomy as ‘the right or condition of self-government’ and the ‘freedom of external control or influence’. It follows that autonomy is a relational factor between two or more actors, parties or entities such as people, governments, computer devices, or intelligent agents. When one actor has an influence over the way a second actor operates, the second’s autonomy is reduced. It should be noted that dependences between actors can sometimes affect both actor’s levels of autonomy. Returning to the example given earlier, many people may remember gaining a lot of autonomy when they are no longer dependent on their parents for transport; but we must also consider the parent – those with children may remember the great sense of relief when their child was no longer dependent on them to be their personal taxi. Here the parent has regained some autonomy as the child’s needs and dependency no longer influences their daily activities.

Margaret Boden, a pioneer in research on artificial life, provides us with a very nice definition of autonomy [15]:

*“Autonomy is not an all-or-nothing property. It has several dimensions, and several gradations. Three aspects of behaviour – or rather, of its control – are crucial. First, the extent to which response to the environment is direct (determined by the present state in the external world) or indirect (mediated by inner mechanisms partly dependent on the creature’s previous history). Second, the extent to which the controlling mechanisms were self-generated rather than externally imposed. And third, the extent to which inner directing mechanisms can be reflected upon, and/or selectively modified in the light of general interests or the particularities of the current problem in its environmental context. An individual’s autonomy is the greater, the more its behaviour is directed by self-generated (and idiosyncratic) inner mechanisms, nicely responsive to the specific problem-situation, yet reflexively modifiable by wider concerns.”*

As pointed out by Hexmoor et al., having a certain degree of autonomy is a defining characteristic of agents [16]. But this begs the question – how much autonomy do we really want our computer systems to have? What’s more, there are obvious advantages to having autonomous-agents that can perform actions on our behalf but does this added convenience come at a price?

These questions are not easy to answer. They depend heavily on the application domain, task at hand, and the ability of the agents. Parasuraman et al. present a ten-point scale (shown in Table I) for possible levels of automation of decision making and action selection [17], which is based on earlier work of Sheridan and Verplank [18]. The levels of automation described easily map to levels of autonomy for agents. The higher up the scale, the more automated (or autonomous) the system becomes. At the highest level the computer acts completely autonomously and ignores the

TABLE I. Levels of Automation of Decision and Action Selection

| | | |
|---|---|---|
| *High* | 10 | The computer decides everything, acts autonomously, ignoring the human |
| | 9 | Informs the human only if it, the computer, decides to |
| | 8 | Informs the human only if asked |
| | 7 | Executes automatically then necessarily informs the human |
| | 6 | Allows the human a restricted time to veto before action execution |
| | 5 | Executes a given suggestion if the human approves |
| | 4 | Suggests one alternative |
| | 3 | Narrows the selection down to a few |
| | 2 | The computer offers a complete set of decision/action alternatives |
| *Low* | 1 | The computer offers no assistance: human must take all the decisions and actions |

 

human user. A little lower down the scale, at level 6, the computer decides which action to take but allows the user a restricted amount of time to veto the action before it occurs. At level 3, the computer determines a suitable set (narrowed down from the full set) of possible actions and allows the user to choose. At the lowest level, the computer does not provide any assistance and the user must make all decisions and perform all actions.

Parasuraman et al. discuss how automation can be applied at varying levels across four different functional stages in computer systems: information acquisition, information analysis, decision selection, and action implementation [17]. At each of these different functional stages, various forms of system autonomy (e.g. along the ten-point scale) are achievable. Parasuraman et al.'s functional stage model can be mapped directly on to agent functionality in an intelligent environment. For example, information acquisition becomes sensing the environmental state, information analysis maps to processing sensed information (finding meaningful data to present to the user or use for reasoning about actions), decision selection is akin to deciding from a set of actions to execute in the environment, and action implementation naturally maps to the execution of actions in the environment. Aside from these, as we are dealing with autonomous-agents and not just automated systems, we also have to account for an extra function: governance of the system (e.g. learning new actions, evaluating performance, and re-learning based on performance). Governance is a vital aspect of intelligent environments; without it an intelligent environment would not be able to adapt to the user or to the changing conditions that is faces.

Parasuraman et al. outline a number of factors a system designer should consider when deciding on what levels of automation or autonomy to use when designing systems. Primarily, Parasuraman et al. suggest designers should consider the human performance consequences of the resulting system: mental workload, situation awareness, complacency, and skill degradation. The fundamental aim of automation and autonomous systems is to reduce or ease the mental workload of people; however, a poorly designed autonomous system may effectively increase the cognitive effort by people, for example when the autonomous functionality is difficult to initiate or requires a tedious amount of data entry [19]. Situation awareness, complacency, and skill degradation are all closely related issues that are underpinned by a user's over-reliance, over-trust, or over-use of high-level automation and high autonomy systems. In these cases, when the computer makes all the decisions or actions on behalf of the user, the user no longer has to actively think about what's happening to cause these decisions/actions, and hence are not always fully aware of the current situation [20]. This could lead to users misunderstanding the computer's actions and ultimately not catching mistakes made by the computer. A similar case may also arise if the user becomes too complacent in trusting the computer – there is a danger that the user will simply assume that the computer is correct and perhaps miss a critical mistake by the computer [21]. Furthermore, over time the user's ability and skills may degrade as the computer takes control of certain decisions and actions [22]. This could cause problems if the computer fails at any point and again may reduce the ability of the user to spot the computer's mistakes.

As further considerations for designing automation and autonomous systems, Parasuraman et al. suggest automation reliability and the cost of decision/action outcomes [17]. Reliability of the automation/autonomous system is big concern as it directly affects a user's trust in the system, which ultimately affects how the system is used (or under-used) and, in extreme cases of unreliability, whether it's used at all [23]. It's generally acknowledged that absolute reliability is not realistically achievable and computers, intelligent agents and machines are expected to fail at some point. Hence, it is necessary to consider the cost of decision/action outcomes of automated and autonomous systems alongside the reliability of the system. For example, if there is a very high cost associated to a wrong decision then this decision should not be made autonomously by a computer unless it also has very high reliability (i.e. the possibility of the cost being incurred is very low).

In terms of intelligent environments, we must consider an even wider set of criteria when reasoning about autonomy because of how deeply embedded into our lives pervasive technology aims to be. We must not only take into consideration technological limitations of the system and personal concerns, satisfaction, and well being of the user but must also adhere to an extensive set of social constraints and norms of the environment in which the technology operates, which are extremely hard for an agent to sense or model. What's more, all of these (technological limitations, personal factors, and social constraints) are likely to vary over time as well as changing between different environments. In light of this, this work encourages the use of an adjustable autonomy mechanism in agents. Firstly, this allows for an agent's level of autonomy to be adapted the current contextual needs and, secondly, if the user is given control if this mechanism, this allows the user to set the agent's autonomy to suit their personal preferences, needs and concerns.

Adjustable autonomy, as described by Bradshaw et al., allows a system to be "governed at a sweet spot between convenience (i.e. being able to delegate every bit of an actor's work to the system) and comfort (i.e. the desire to not delegate to the system what it can't be trusted to perform adequately)" [24]. Bradshaw et al. describe a general method for adjusting the autonomy of agents that works by: adjusting permissions – allowing and disallowing certain actions in the environment; changing obligations – assigning and withholding tasks to and from the agent; restricting possible actions – for example by restricting resources to the agent; and adjusting the capabilities of the agent – changing the functionality of the agent [25].

Various examples of successful adjustable autonomy systems can be found in recent research in the fields of robotics and artificial intelligence [26, 27]. What's more, researchers in the fields of pervasive computing and ambient intelligence have started to recognise the advantages accruing from an adjustable autonomy approach to managing systems in intelligent environments. For example, researchers at the University of British Columbia, Vancouver, have identified that a central issue in the operational effectiveness of intelligent buildings is the issue of whether 'intelligence' is derived either implicitly or explicitly from the occupants [13]. Also, researchers at Herriot-Watt University, working on the EU funded PERSIST project, noted that the use of exclusive end-user driven methods were not popular with users owing to the high cognitive load placed on the users and that agent assistance was an advantage. In the case of PERSIST the system runs autonomously requesting assistance from the user when it encounters uncertainty in its decision making process [12].

With the extra dynamic of adjustable autonomy in agents, one can allow for control of the system or specific tasks to flow freely between agents and users; hence, in intelligent environments, the user is able to choose how much they wish to control the system and how much trust they are happy to place in the system to operate autonomously on their behalf. The next section discusses how adjustable autonomy can be applied to intelligent environments.

## III. Adjustable Autonomy in Intelligent Environments

As previously discussed in Section II, the functional stages of intelligent environment management can be described at a high level as: sensing, data processing, decision selection, action execution, and governance. Governance incorporates all the learning and adaptive functionality of the intelligent environment, which in turn provides the foundations for the data processing and decision selection stages – mappings of environmental state (or other trigger) to corresponding actions must be learnt and adapted[1]. As governance is such an integral functionality on which the actual (autonomous) intelligence of the environment so heavily depends, this work focuses on how agent autonomy can be altered in the task of governing the system; more specifically, how behaviour rules for achieving automation in the environment (e.g. IF *the room is occupied,* THEN *turn on the lights)* are input into the system and how these rules can be adapted and changed over time.

The Adjustable Autonomy Intelligent Environment (AAIE) architecture model has been designed to enable adjustable autonomy with respect to governance at four different levels of autonomy:

1) *Full autonomy*: the agent learns from the user's behaviour, automatically creates rules accordingly and adapts them over time as the agent deems it necessary.

2) *High autonomy*: the agent learns from the user's behaviour and generates rules accordingly but the rules can only become active in the system when they have been confirmed by the user. Similarly, changes to rules must be first confirmed by the user. At the confirmation stage, the user is presented with the opportunity to accept, reject or edit rules.

3) *Low autonomy*: the user programs the rules and can later modify them using a GUI. The user is assisted by the agent presenting suggestions for rules upon request.

4) *No autonomy:* the user programs and modifies the rules using a GUI with no assistance from the agent.

In a perfect system, one might imagine a much more continuous adjustment of autonomy, perhaps along a sliding scale similar to a volume control switch or dial that allowed the user to tweak the level of autonomy to their liking. However, in a practical sense this is extremely hard to implement as the tasks carried out by agents generally have a discrete number of steps or discrete stages at which a user can take control or have influence over. Hence, the AAIE model should be seen as a step towards a 'more-perfect' adjustable autonomy system in which continuous adjustment is possible.

The AAIE model takes the form of an event driven multi-agent architecture and is based around the framework of the University of Essex iSpace [4]. The iSpace, shown in Figure 1, is a purpose built test bed for intelligent environment and ambient intelligence systems. As well as everything one might expect to find in any other two-bed apartment, the iSpace is also equipped with a multitude of networked sensors and actuators, e.g. internal and external temperature and lighting sensors, real-time location tracking, computer-controllable heating and lighting, and electronically controlled curtains and doors. In Figure 2, the overall architecture of the AAIE model can be seen. In the system there is: the physical environment, which contains numerous devices, sensors and actuators; a Context Agent (CA); an Acting Agent (AA); an Interface Agent (IA); and an Adjustable-autonomy Behaviour-Based Agent (ABBA). The CA monitors the current state of the intelligent environment by communicating with the physical environment via UPnP, listening for events and maintaining up-to-date sensor readings. When an event is detected by the CA (i.e. there is a significant change in the state of the environment) the new environment state and event information are passed to ABBA and used to decide on actions to perform in the environment or learn new rules. The AA drives actuators and devices in the physical environment as instructed by the Coordinator of ABBA. The IA provides an interface between the system and the user; it allows this user to directly control devices in the physical environment using a

Figure 1. The Essex iSpace

[1] This is assuming an agent is handling these tasks and not the user.

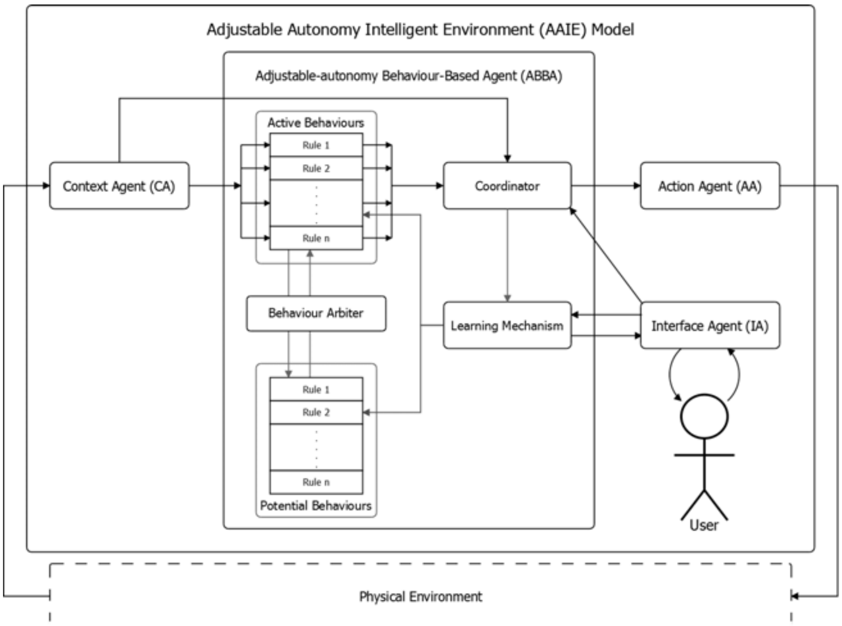


Figure 2. The Adjustable Autonomy Intelligent Environment (AAIE) architecture model.

GUI interface and use the rule generating tools/procedures made available by ABBA.

ABBA provides the controlling and learning ability in the system. It is inspired by another management system for intelligent environments, the Incremental Synchronous Learning (ISL) agent developed by Hagras et al. [18]. The ABBA architecture takes the general form of a behaviour-based architecture, as pioneered by Brooks at MIT [29]. In such architectures, a number of agent behaviours (known as behaviour rules in our system) run in parallel and a controller (named coordinator in our system) is employed to coordinate the behaviours or their given outputs into one single output to achieve the desired agent functionality. As seen in Figure 2, the ABBA architecture contains the following components: the coordinator, the learning component, two sets of behaviour rules and the behaviour arbiter.

In the behaviour rule sets, a behaviour rule take form of an 'IF *state* THEN *action'* mapping, which we refer to as a behaviour rule and describe a set of automated behaviours to be carried out by the agent; for example, IF *the user is in the Living Room,* THEN *turn on the television*. Behaviour rules may be either generated by ABBA or programmed by the end-user. Each rule is also assigned with a confidence level when it is created: a value between 0 and 1, based on whether it is was created by ABBA or the end-user and in accordance with the selected autonomy level (this is explained in more detail later). A rule can only have an effect on the environment if it is active and can only be active if it has a high enough confidence level. Rules with a low confidence can only be potential behaviour rules and cannot effect the environment. All behaviour rules are visible to all components of the agent. The behaviour arbiter component regulates the behaviour rules. The confidence of all rules slowly degrades overtime, and if an active rule's confidence level drops below a certain threshold then it is dropped down into the potential set and if a potential rule's confidence level drops below a very low threshold (zero for example) then it is deleted. This confidence degradation reduces the chance that the agent's memory will become full. To adhere to the "User is King" clause, created by Callaghan et al. [30], a user programmed rule will always have a confidence level of 1, which does not degrade over time, and a user may always program behaviour rules regardless of level of autonomy.

When an event is detected by the CA (i.e. there is a significant change in the state of the environment or a user action) the new environment state and event information are used as an input to the active behaviour rules and also passed to the coordinator. The active behaviour rules, running in parallel, then produce an output based on the input data if possible. The coordinator regulates/merges the output of all the active behaviours into one single output so that each behaviour rule effects the external environment to an appropriate degree. Additionally, when an active behaviour rule affects the environment the coordinator increases the level of confidence of that rule, where the amount increased depends on the degree the rule is effecting the environment. Thus, the more a rule is used and the more it effects the environment, the less chance it will have of dropping into the potential set and ceasing to be active.

If no behaviour rule produces an output for any given event, then the event information is sent to the ABBA's learning component. If this event is a user action, then the learning component uses the event information to generate a new rule, which is assigned a very low level of confidence (for example 0.1) and is placed in to the potential rule set. If the same user action is observed subsequent times, the confidence of the newly generated rule is raised by a certain amount (incremented by 0.1 for example). A clustering algorithm and genetic algorithm (GA) is used to find similar behaviour rules (i.e. rules derived from similar user behaviour) and merge them together to find (more optimal) new behaviour rules. If a more optimal set of behaviour rules can be found in this way, the confidence of the new rules is increased over that of the original (merged) rules, and if the confidence level is then high enough, the rule may become active.

To achieve different levels of autonomy in ABBA, we can alter how behaviour rules are generated and become active in the system. Firstly, a confidence level, between 0 and 1, is assigned to each behaviour rule, and the agent is restricted so that a behaviour rule can only be active (and affect the environment) if the confidence level is high enough. By way of example, let us say that for a rule to become active in the system it requires a confidence level of 0.9. Then at full autonomy, the agent can learn new rules automatically and can assign anything up to a confidence level of 1 to rules (i.e. no direct interaction between the agent and the user is required). However, at no autonomy, the agent is restricted so that it cannot learn new rules and cannot assign confidence to existing rules; hence, it is the sole responsibility of the user to manage the system and the agent has no affect. For high autonomy, the agent is restricted so that it can generate new rules but can only assign a confidence level up to a cap of 0.75 (for example). Once this cap is hit, the agent must then communicate with the user to confirm the rule in order to attain the extra 0.15 confidence required for the rule to become active. Here, since the system requires an input from the user, we can say that the agent is no longer fully-autonomous. For a low autonomy level, the agent is allowed to create new rules but not to assign confidence to the rules; instead the agent uses its known rules as an experience bank to form suggestions and aid the user in programming their desired rules. Here, the user has the majority of influence over rule creation but can request the agent to provide some input, giving the agent a low level of autonomy.

By enabling adjustable autonomy in an intelligent environment we can allow the user to change how much control they give to agents and how much control they wish to maintain themselves depending on their attitudes towards the system, devices and agents. The next section discusses the major factors and concerns of users that determine their preferences and selection of different autonomy levels.

## IV. Concerns and Attributes Affecting Autonomy Selection

An online survey was recently conducted, which aimed to investigate people’s opinions of the use of autonomy in intelligent environments (using the example of a smart home) [14]. As a follow up to the online survey, a working adjustable autonomy intelligent environment (using the AAIE model) was implemented and a series of user trials were conducted, which aimed to gain deeper insights into the reasoning behind people’s attitudes of different levels of autonomy and explore how using adjustable autonomy can change people’s opinions of intelligent environments [14]. The results showed that people have many different concerns when it comes to ambient intelligent systems and their attitudes towards autonomous-agents are highly individual and differ greatly between people. Furthermore, the results strongly indicate that different people may prefer different levels of autonomy in different situations and for different sub-systems of an intelligent environment, plus their views may drift over time (e.g. as they learn more about consequences of using the technology). The major findings from these studies highlight the main determining factors and concerns of users that have with respect to their attitudes and preference of autonomy in intelligent environments. These major findings are discussed in summary in this paper; for a more in depth discussion of both of these studies please refer to [14].

As discussed in Section II, Parasuraman et al outline a number of key considerations for producing automated systems [17]: mental workload, situation awareness, complacency, skill degradation, reliability of the automation, and the cost of decision/action outcomes. All of these issues were voiced in one way or another by the participants of the online survey and user trials whilst discussing their concerns with autonomous-agents.

*Mental Workload:* the vast majority of participants could recognise the benefit of using autonomous-agents and could see how it potentially could reduce the cognitive load placed on them in governing the intelligent environment’s rule-based system. Many identified the full and high autonomy levels as being especially useful for people with a “busy life style” as it would save them from having to “come up with a suitable list of rules”, program these rules and maintain them overtime as the would with the low or no levels of autonomy. Conversely, many were concerned that the agent might perform badly or simply be unable to learn from their behaviours and daily routines, and hence felt that allowing the agent full autonomy may in fact mean they repeatedly have to correct the agent’s mistakes. What’s more, with the high level of autonomy, some felt that they may become annoyed by repeatedly having to deal with the agent’s suggestions for rules.

*Situation awareness:* many participants felt that having an agent with full autonomy would reduce their awareness of what was happening in their homes and didn’t like the idea that an agent could program rules. Some said they would “find it quite scary” if things start to happen in their homes unexpectedly. However, there were a small number of participants that felt having the help of an agent would in fact enhance their situational awareness. For example, some thought it may be a good way to monitor the habits of themselves and their families; for example, one participant thought that allowing the agent to learn autonomously may be a good way to “keep an eye on how much TV the kids are watching”.

*Complacency and skill degradation:* very few participants expressed concerns about becoming overly complacent with autonomous-agents or losing personal skills or abilities. This is perhaps because the application domain of someone’s home is somewhere that most would like to feel complacent or perhaps due to their opinion that the agent wouldn’t perform well in real life. However, some did have concerns with becoming overly dependent on the technology and worried that it may affect their personal health by making them become overly lazy.

*Reliability:* This was also one of the biggest concerns expressed by the participants. Many doubted the ability of the agent to recognise their more complex behavioural patterns, which depend more on personal feelings and mood. Hence for activities such as control of entertainment and media devices it was found that many would prefer to set an agent to low or no autonomy; however, for devices such as lighting and air conditioning quite a few said they would rather use full or high autonomy as they felt the usage of these devices is more ‘routine’ depending more simply on the time of day and/or state of the environment, and so decided it would be quite easy (and more convenient) for the agent to deal with. It should also be noted that many said they would give the agent a chance to prove itself worthy in handling these more complex tasks, and if it could then they might assign it more autonomy. The issue of reliability is closely related to the concept of *trust* in the system, which is very fundamental concern in user-centric systems [31]. In our day-to-day lives, trust plays a heavy role in how we decide on our personal relationships with others. Usually we will seek closer relationships with those who we trust more and try to keep distance with those we deem less trustworthy. The same follows for how we develop relationships with computer systems in our lives. For us to continue using the calendar service on our smart phones, for example, we must place a lot of trust in it to remind us of appointments in an appropriate and timely manner; when it fails to do so (i.e. the service is unreliable) our trust is broken and the user-system relationship breaks down.

*Cost of decision/action outcomes:* in the context of smart homes, the cost of failures was found to be concern which mainly related to user annoyance. People feared that the agent would repeatedly make mistakes at full autonomy, which they would have to keep correcting, or would keep making unnecessary suggestions at the high autonomy level, which they would have to deal with. In all cases where these fears arose, they seemed to be driven by people’s previous (bad) experience with similar technologies and assistive systems. Contrary to this, a smaller number of people expressed concerns over their own ability to manage such a complex system, and were worried about the costs of their own mistakes somehow ‘breaking the system’; hence, they felt that it would be better to allow a more-autonomous agent to govern the system on their behalf. It was found that the perceived cost of (erroneous) decisions and actions is heavily dependent on the *type of device* being controlled. For example, in the user trials it was found that many people wouldn’t mind giving control of air conditioning and lights to an agent as the usage of these devices is generally not too delicate or sensitive, hence it doesn’t matter too much if an agent makes mistakes. With the curtains, however, many did not want to give too much autonomy to the agent as they would be worried that the curtains could open unexpectedly at inappropriate times.

As can be seen, the considerations for designing automated systems, as outlined by Parasuraman et al., do apply for an autonomous-agent system such as smart homes. However, as smart homes and other intelligent environments are very user centric systems, a much deeper set of user concerns need to be considered. For example, the major issue of trust in the agents and the ways in which different type of device being controlled affects the user's opinions has already been discussed. Further concerns raised by the study participants were: maintaining control, privacy, the current social context, and personal feelings towards technology.

*Control:* alongside reliability of the agent, feeling in control over what happens in the environment was one of the biggest concerns of the participants. Many expressed concerns relating back to awareness of what’s happening in the system when using agents with full autonomy. However, it was found that people’s concerns of control were greatly reduced if the agent operates at the high level of autonomy instead of the full. The participants preferred the high level as the agent “put it in writing first” and they maintain direct control over the agent by having the option to confirm the agent-derived rules before they become active.

*Privacy:* people’s homes are perhaps the most intimate and personal environments in our lives; it follows that people would want most if not all of their activities in the home to be kept private [32]. It was found that many people’s privacy concerns were eased somewhat with the inclusion of the adjustable autonomy mechanism as this gave the user the ability to stop agents from monitoring parts of the environment (by switching the agent to no autonomy). Although, there were some that still feared others having external access to their personal data even with agents operating with no autonomy; for example, one experiment participant said they would be worried if the government could access the data and how it might be used, or rather misused as in Orwell’s Nineteen Eighty-Four [3].

*Social context:* the home is a very social place and many participants felt there would be a need to dial down agent autonomy (to turn off the learning ability and acting ability of the agent) when people came to visit as they felt that many of their personal behaviour rules for automation might not be appropriate, and they didn’t want agents to learn from visitors’ actions in the environment.

*Technophobia/Technophilia:* throughout the experiments people expressed a variety of attitudes towards technology. Some people were found to have a general fear or dislike of technology (especially in the pervasive sense) while others are extremely excited by anything that is high-tech. One extreme case of technophobia was one participant who felt that having computers that take control of our living space is “taking our humanity” and they thought that human beings should strive to be as independent from machines as possible. On the opposite end of the scale, another participant had a very strong appreciation for high-tech gadgets and said they would love to live in a smart home so they could show off to their friends.

These user concerns build on those outlined for the design of automation by Parasuraman et al. As can be seen, there are heavily contrasting viewpoints for each of the concerns or issues that can affect a user's preference towards levels of autonomy. With such an array of mixed opinions, which often change over time, these concerns in fact go far beyond the perception and understanding of the system designer. Hence there is a real need to employ adjustable autonomy in intelligent environments to allow a user to explore the trade-off between convenience of higher autonomy and control of lower autonomy and to find their own personal sweet-spot, and alter at a later date if they see fit

## V. CONCLUSION

There is a long-standing debate over the use of autonomous-agents in intelligent environments. While many believe that research should focus on developing end-user driven systems, seeking to empower the user, many others maintain that intelligent environments should be autonomous-agent driven, minimising the work and effort required from the user. Both of these approaches have their distinct advantages, but they are not suitable for all. This work explores how it is possible to make intelligent environments more dynamic and personalisable by equipping them with adjustable autonomy, which allows the user to alter the amount of influence autonomous-agents have over managing their intelligent environment. A recent online survey and user study were conducted to gauge people's opinions on the use of autonomy in intelligent environments, for which the Adjustable Autonomy Intelligent Environment (AAIE) model was implemented as an experimental system in the University of Essex iSpace. This paper has outlined a number of key factors and concerns of users that determine their personal preferences towards different levels of autonomy. Given the shear diversity between these variables, equipping intelligent environments with an adjustable autonomy system is extremely useful as it allows the user to find their personal and very individual sweet-spot between control and convenience, rather than the decision being left to the system designer.

## REFERENCES


[1] Castelfranchi, C., Falcone, R., "From Automaticity to Autonomy: The Frontier of Artificial Agents", in Agent Autonomy, Springer, pp.103-136, (2003).

[2] Bentham, J., "Panopticon Writings", Pub: Verso, ISBN-13: 978-1859840832, (1995)

[3] Orwell, G., "Nineteen Eighty-Four", Penguin Books, (1949).

[4] Callaghan, V., Clarke, G.S., and Chin, S.J.Y., "Some Socio-Technical Aspects Of Intelligent Buildings and Pervasive Computing Research", Intelligent Buildings International Journal, Earthscan, 1:1, (2008).

[5] Humble, J., Crabtree, A., Hemmings, T., Åkesson, K.P., Koleva, B., Rodden, T., Hansson, P., "'Playing with the Bits', User-Configuration of Ubiquitous Domestic Environments", in Proceedings of UbiComp 2003, Springer-Verlag, pp.256-263, (2003).

[6] Gajos, K., Fox, H., Shrobe, H., "End user empowerment in human centered pervasive computing", in Proceedings of Pervasive 2002, pp.1-7, (2002).

[7] Chin, J., Callaghan, V., Clarke, G., "An End User Tool for Customising Personal Spaces in Ubiquitous Computing Environments", in Lecture Notes in Computer Science: Ubiquitous Intelligence and Computing, Springer-Verlag, pp.1080-1089, (2006).

[8] Callaghan V., Colley M., Hagras H., Chin J., Doctor F., Clarke G., "Programming iSpaces: A Tale of Two Paradigms", in Intelligent Spaces: The Application of Pervasive ICT , Springer-Verlag, (2005).

[9] Hagras, H., Callaghan, V., Colley, M., Clarke, G., Pounds-Cornish, A., Duman, H., "Creating an Ambient-Intelligence Environment Using Embedded Agents", in IEEE Intelligent Systems, Vol.19, No.6, pp.12-20, (2004).

[10] Mozer, M.C., "Lessons from an adaptive home", in Smart Environments: Technology, Protocals and Applications, Wiley, pp.273-298, (2005).

[11] Ball, M., Callaghan, V., "Perceptions of Autonomy: A Survey of User Opinions Towards Autonomy in Intelligent Environments", in Proceedings of the 7th International Conference on Intelligent Environments 2011, Nottingham, UK, IEEE Xplore (2011).

[12] Gallacher, S., Taylor, N., Williams, H., Papadopoulou, E., "Giving the User Explicit Control over Implicit Personalisation", in Proceedings of the PERSIST Workshop on Intelligent Pervasive Environments : AISB 09, Edinburgh, UK, (2009).

[13] Cole, R.J., Brown, Z., "Reconciling human and automated intelligence in the provision of occupant comfort", in Intelligent Buildings International, Volume 1, pp. 39-55(17), (2009).

[14] Ball, M., Callaghan, V., "Managing Control, Convenience and Autonomy in Special Issue on Agent-Based Approaches to Ambient Intelligence, AISE book series, IOS Press, (2012).

[15] Boden, M., "Autonomy and Artificiality", in The Philosophy of Artificial Life, Oxford University Press, pp.95-108, (1996).

[16] Hexmoor, H., Castelfranchi, C., Falcone, R., "A Prospectus on Agent Autonomy", in Agent Autonomy, Springer, (2003).

[17] Parasuraman, R., Sheridan, T., Wickens, C., "A Model for Types and Levels of Human Interaction with Automation", in IEEE Transactions On Systems, Man, And Cybernetics, Vol. 30, No. 3, (2000).

[18] Sheridan, T., Verplank, W., "Human and Computer Control of Undersea Teleoperators," MIT Man-Machine Systems Laboratory, Cambridge, MA, Tech. Rep., (1978).

[19] Kirlik, A., "Modeling strategic behavior in human-automation interaction: Why "aid" can (and should) go unused," in Human Factors, vol. 35, pp. 221–242, (1993).

[20] Sarter, N., Woods, D., "'Strong, silent, and out-of-the-loop': Properties of advanced (cockpit) automation and their impact on human-automation interaction," Cognitive Systems Engineering Laboratory, Ohio State University, Technical Report CSEL 95-TR-01, (1995).

[21] Parasuraman, R., Molloy, R., Singh, I., "Performance consequences of automation-induced 'complacency'," in Int. J. Aviation Psychology, vol. 3, pp. 1–23, (1993).

[22] Rose, A., "Acquisition and retention of skills," in Application of Human Performance Models to System Design, G. McMillan, Ed. New York: Plenum, (1989).

[23] Parasuraman, R., Riley, V., "Humans and automation: Use, misuse, disuse, abuse," in Human Factors, vol. 39, pp. 230–253, (1997).

[24] Bradshaw, J. M., Sierhuis, M., Acquisti, A., Feltovich, P., Hoffman, R., Jeffers, R., Prescott, D., Suri, N., Uszok, A., & Van Hoof, R., "Adjustable autonomy and human-agent teamwork in practice: An interim report on space applications", in Agent Autonomy H. Hexmoor, R. Falcone, & C. Castelfranchi (Eds.), pp.243–280, (2003).

[25] Bradshaw, J.M., Feltovich P.J., Jung, H., Kulkarni, S., Taysom, W., Uszok A., "Dimensions of adjustable autonomy and mixed-initiative interaction", in Agents and Computational Autonomy: Potential, Risks, and Solutions, Springer-Verlag, pp.17–39, (2004).

[26] Brookshire, J., Singh, S., Simmons, R., "Preliminary results in sliding autonomy for coordinated teams", in Interaction between humans and autonomous systems over extended operations, AAAI Spring Symposium, pp.127–132, (2004).

[27] Bruemmer, D., Dudenhoeffer, D., Marble, J., "Dynamic-Autonomy for Urban Search and Rescue", AAAI Technical Report WS-02-18, (2002).

[28] Hagras, H., Callaghan, V., Clarke, G., Colley, M., Pounds-Cornish, A., Holmes, A., Duman, H., "Incremental Synchronous Learning for Embedded-Agents Operating in Ubiquitous Computing Environments", in Soft Computing Agents: A New Perspective for Dynamic Information Systems, IOS Press, (2002).

[29] Brooks, R., "Intelligence Without Representation", in Artificial Intelligence, 47:139-159, (1991).

[30] Callaghan V., Colley M., Hagras H., Chin J., Doctor F., Clarke G., "Programming iSpaces: A Tale of Two Paradigms", in Intelligent Spaces: The Application of Pervasive ICT , Springer-Verlag, (2005).


 

[31] Trcek, D., ”Trust Management in the Pervasive Computing Era,” in Security & Privacy, IEEE , vol.9, no.4, pp.52-55, (2011).

[32] Spiekermann, S., “Perceived Control: Scales for Privacy in Ubiquitous Computing”, in Proceedings of the 10th International Conference on User Modeling, (2005).