FROM PASSIVE TO PROACTIVE DESIGN ELEMENTS

Incorporating curious agents into intelligent rooms

MARY LOU MAHER
University of Sydney, Australia

and

KATHRYN MERRICK
University of Sydney and National ICT Australia, Australia

and

ROB SAUNDERS
University of Sydney, Australia

Abstract. Agent technology has been used as an organising mechanism for software systems that focus on modularity and autonomy. This paper presents two applications that explore the potential of combining agent technologies with physical building design elements to change the nature of the built environment from a passive space to one that proactively engages with its inhabitants. We focus on how these curious places sense the state of the environment and the activities of the humans in the environment and *enhance* the human experience, thus going beyond the concept of *supporting* human activities in traditional approaches to intelligent rooms.

1. Introduction

Agent technology provides a model for an autonomous reasoning entity that senses its environment and determines a response as an action on that environment. Agents have been used as an organising mechanism for software systems that focus on modularity and autonomy (Wooldridge and Jennings 1995). In this paper we go beyond software systems of agents and look at how agent technology can sense and effect physical building design elements, changing the nature of the built environment from a passive

A Dong, A Vande Moere & JS Gero (eds), CAADFutures'07, 447-460.

environment to one that proactively engages the environment and its inhabitants. We focus on how curious places that sense both the state of the environment and the activities of the humans in the environment can *enhance* the human experience, going beyond the concept of *supporting* human activities in traditional approaches to intelligent rooms.

Our basic approach builds on previous work in computational models of curiosity, novelty (Saunders 2001), motivation and learning (Merrick and Maher 2006) in agents. Artificial curiosity is modelled on Berlyne's psychological theory, which states that the most curious sensations are those that are moderately stimulating in terms of their similarity to previously encountered stimuli (Berlyne 1960).

2. Approaches to Intelligent Rooms

The concept of an intelligent room is a physical space for living or working that includes embedded computational power to facilitate or augment the actions of users of the environment as they perform their daily tasks. Intelligent rooms monitor the activities that take place within them using sensors and respond to sensations using effectors in order to exhibit intelligent behaviour and assist users. Sensors may include devices such as motion detectors or pressure pads while effectors may include devices such as lights, projectors or doors.

Research in intelligent rooms can be regarded as a sub-field of ubiquitous computing which aims to integrate computers seamlessly into everyday living. Brooks (1997) and Coen (1998) have argued that intelligent rooms should:

- Adapt to and be useful for ordinary everyday activities,
- Assist the user without requiring the user to attend to them,
- Have a high degree of interactivity,
- Be able to understand the context in which people are trying to use them and behave appropriately. (Kulkarni 2002)

An intelligent room is essentially, as Kulkarni (2002) suggests, an immobile robot. However, the design requirements of an intelligent room differ from those of normal robots in that they are oriented towards maintaining an internal space rather than exploring or manipulating an external environment.

Current agent based approaches to intelligent room design include MIT's intelligent room prototype e21, which facilitates activities via a system called ReBa (Hanssens et al 2002). ReBa is a context handling system that observes a user's actions via the reports of other agents connected to sensors in the room's multi-agent society and uses them to build a higher-level representation of the user's activity. Each activity – such as watching a movie or giving a presentation – has an associated software agent, called a behaviour agent. The behaviour agent responds to a user action and performs a reaction, such as turning on the lights when a user enters the room. Behaviours can form layers based on the order of user actions, acknowledging differences in context such as showing a presentation in a lecture setting versus showing one in an informal meeting. Although ReBa can infer context in this way, it cannot adapt to patterns of usage: to create an entirely new context, ReBa's behaviour agents must be reprogrammed to recognize the user's actions and take an appropriate reaction. It does not self-adapt to new usage patterns.

Other researchers have taken approaches to designing intelligent rooms that are not explicitly agent-based. Both the University of Illinois' Gaia Project (Roman et al 2002) and Stanford University's Interactive Workspace Project (Johanson et al 2002) have taken an operating systems approach, developing Active Spaces and Interactive Workspaces respectively, which focus on the role of the physical space as a platform for running applications rather than as a proactive facilitator: in these systems, actions are triggered by the user and behaviours are pre-programmed by the applications developer. Gaia's context service provides the tools for applications developers to create agent-based facilitating applications so the overall model is reactive rather than adaptive.

Our approach to an intelligent room uses an adaptive agent model (Maher et al 2006). An agent is a system that perceives its environment through sensors and uses some characteristic reasoning process to generate actions using effectors. Agent models correspond well to intelligent rooms as both are described as having sensors for monitoring their environment and effectors for making changes to the environment. A variety of agent models have been developed with different characteristic reasoning processes for mapping sensor input to effector output. These range from simple rule-based reactive agents to complex cognitive agents that try to maintain and reason about an internal model of the world using planning or machine learning algorithms. This raises the question of what kind of agent model is most suitable as a basis for an intelligent room.

In this paper we present two applications that contribute to the development of models for proactive, intelligent rooms as a new kind of building design element. The first is a curious information display that observes the movement of people in the room and tries to attract their attention by learning different information display combinations. The second is a curious researcher that observes the content of the research presentations given in the room and generates its own research presentations that are interesting to the human researchers. From these applications we can draw some conclusions on the type of agent models that can become part of building design elements to transform the built environment from a passive space to a proactive place.

3. Curious Information Display

A curious information display augments physical places by attracting the interest of observers and selecting information to display by being curious and learning about the structure and content of the information it presents. Traditional information displays such as posters and billboards present a fixed image to observers. Recently, digital displays have become a popular means of presenting information. Digital displays allow the amount of information presented to be increased, often by attaching the display to a computer that changes the contents of the display automatically. The full power of a digital display has not yet been realised, with the most common scenario using a database of images and displaying images at random or in a predefined order. In this section we describe agent models that provide the infrastructure for a proactive information display, intended to attract and respond to the people in the room.

In scenarios where the display of digital information is more familiar, such as web-browsers, novel interaction algorithms have been developed to automatically personalise the digital space (Dieterich et al 1993). Similarly, intelligent tutoring systems use artificial intelligence algorithms to tailor learning material to the individual needs of students (Graesser et al 2001). Large digital information displays in public spaces have the same capacity for the use of novel techniques to improve the usefulness of the displays. However the public, multi-user nature of these displays calls for new algorithms to improve the ability of such displays to impart information.

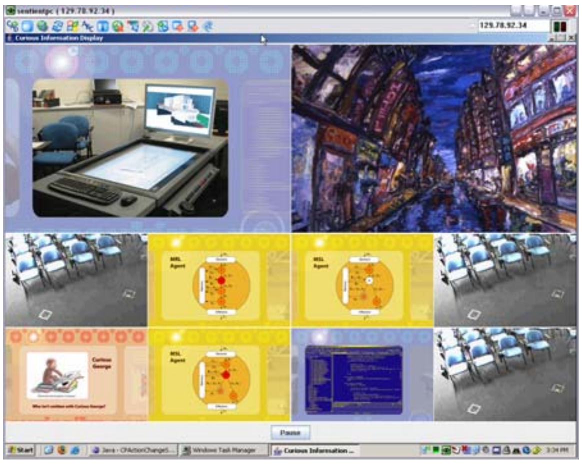

Figure 1. The curious information display.

This section introduces two models of curious information displays that display information about design, computing, agents, and curiosity. Information on these themes is obtained from a research image database, the world-wide-web and live webcam images. The display is located in a research seminar room and uses a motivated reinforcement learning agent model (Merrick and Maher 2006) to detect and learn interesting events.

Our curious information display comprises a matrix of displayed information items (IIs). The source of the II may be a definition or an image from a database, an image from the web, or video from a webcam. Each item can be displayed in a 1x1, 2x2 or 4x4 cell as shown in Figure 1. Each cell is referenced as a leaf node and contains an II from one of the three sources based on one of the keywords design, computing, curiosity, or agent.

The curious information display agent is a motivated reinforcement learning agent. Its reasoning process is decomposed into four sub-processes: sensation, motivation, learning and activation, illustrated in Figure 2.

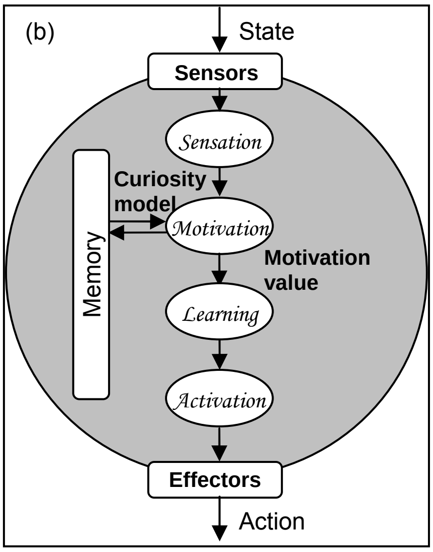

Figure 2. Motivated reinforcement learning agent model used to control the curious information display.

The sensation process transforms raw data from the agent's sensors into structures to facilitate further reasoning. These include the observed state of the environment and the change or 'event' between the current and previous observed states. The motivation process reasons about the current observed state, events, and/or a representation of the set of all observed states encountered so far, to produce a motivation value which is used as an intrinsic reward signal for the learning process. The learning process performs a reinforcement learning update to incorporate the previous observed state, action and current rewards into a policy defining how the agent should act. Finally, the activation process selects an action to perform from the learned policy.

The role of the motivation process in motivated reinforcement learning is to provide an intrinsic reward signal to direct the learning process. Firstly, events and observed states are received from the sensation process. Next, a focus of attention structure is used to distinguish between different input stimuli. One or more characteristic motivation functions are then used to compute motivation values. These motivation values are combined using a reward function which computes a single intrinsic reward signal. This value and the observed state are then passed to the reinforcement learning process.

To create displays that are both interesting to their viewers and also interested in the structure and content of the information they display, we modify the Saunders (2001) model of interest to create an intrinsic motivation signal for motivated reinforcement learning. This model first computes the novelty of a stimulus from the environment using an Habituated Self-Organising Map (HSOM) as a focus of attention mechanism. The characteristic motivation functions are Stanley's model of habituation for computing novelty (Stanley 1976) and the Wundt curve for computing interest (Wundt 1910). The computed value of the interest function is used directly as the reward signal.

In this application, we replace the SOM component of the HSOM with K-means clustering as the attention focus mechanism. In environments such as the curious information display application where sequential stimuli share a large number of common features, SOMs can be dragged towards a corner of the state space when the same neuron is repeatedly selected as the winner. In K-Means clustering, a single neuron can move without deforming the entire network, making it a more appropriate attention focus mechanism in this application.

While the motivated reinforcement learning agent model is designed to be independent of any specific reinforcement learning algorithm, certain classes of reinforcement learning are more appropriate in intelligent room applications where learning is continuous rather than episodic. Temporal difference reinforcement learning algorithms such as SARSA (Rummery and Niranjan 1994) and Q-learning (Watkins and Dayan 1992) are the most appropriate in these settings as they do not require a model of the environment from which to learn and learning occurs after each action that is performed by the agent. As the curious information display has a large problem space we use Q-learning combined with neural network function approximation in the learning process in this application.

We implemented two types of curious information display using different aspects of our model of motivation that offer different capabilities. The first uses events to trigger intrinsic motivation while the second uses observed states.

3.1. A CURIOUS INFORMATION DISPLAY USING INTERESTING EVENTS

This curious information display extends static information displays by reasoning about the changes it can make in the structure and content of the IIs it displays and finding interesting patterns of behaviour to modify the structure and content. This aim of this type of display is to achieve sequences of actions that make interesting changes to the structure and content of information items being displayed. To facilitate this, the motivation process reasons about events in order to identify interesting changes in the display.

One weakness of this agent from a visual perspective is its tendency to favour simple behaviours of only one or two actions. The simplest technique for repeating most of the events in the environment is to cause the event, undo it, then repeat it. This 'shortest path' is naturally favoured by reinforcement learning. This phenomenon is a result of a state space with a moderate level of structure and complexity. There is enough complexity to continually stimulate the agent's motivation process to produce high reward and focus learning on new two-step changes, but not enough structure to motivate the emergence of more complex behaviours as has previously been possible in MRL agents in other environments (Merrick and Maher 2006). Further work is required to understand the impact of the state space structure on learning.

The key characteristic of this type of display is its ability to change rapidly between different configurations and different information content. This is because the MRL agent controlling the display is reasoning about events or changes in the display and is thus motivated to continue to change the layout and content of the display either to focus on an interesting change or to search for new changes that might be interesting. This type of curious information display could be useful as an ambient display device for information in pictorial or diagrammatic form which can be understood at a glance, rather than requiring reading. Changes in the display are eye-catching and the movement between different displays holds the viewers' attention by displaying related information.

3.2. A CURIOUS INFORMATION DISPLAY USING INTERESTING OBSERVED STATES

This curious information display extends static information displays by reasoning about the structure and content of the information items it displays and learning to maintain interesting displays. To achieve this, the motivation process reasons about observed states to identify interesting display states.

Observations show that this display appears to react much more quickly to high reward than the previous display. When high reward is encountered the agent freezes the display within an action or two and focuses on that configuration until reward is reduced and boredom triggers exploration.

In contrast to the previous display, this type of curious information display tends to maintain specific configurations of the display for longer periods of time. This type of curious information display could be useful for both diagrammatic and textual information which requires more time for the viewer to understand. Because this display naturally changes more slowly, sudden changes caused by change in the reward signal are highly noticeable.

4. Curious Research Space

Here we introduce the infrastructure for curious research spaces, which extend the intelligent room with agent technologies that actively contribute to research activities. In a curious research space, agents monitor digital research data produced by humans, such as documents and presentations. They extract keywords to model human data and identify relevant new data from documents on the world-wide-web, which is compared to the human research data model using a computational model of interest to select data to be presented. Interesting new data is then automatically formatted into a slide show presentation be presented at research group meetings.

Much research effort has been applied to the development of web-search engines to extract data from the world-wide-web given keywords, resulting in products such as the Google (www.google.com) and Yahoo! (www.yahoo.com) search engines and web-search clustering tools such as the Carrot2 clustering engine (http://demo.carrot2.org/demo-stable/main). Our curious research space makes use of these tools but extends existing technologies by integrating them into a physical environment. The keywords which trigger a search are extracted from research data generated by human activities in a physical space and the large amounts of data which result from a search are reduced and focused using computational models of interest to value search results. The curious research space also contributes to the field of automatic document generation, by using a grammar for producing slide shows automatically.

The agent environment is defined by the people and the research data presented in a single room. The room includes a computer with a database into which human researchers can enter the URL of research data they present at meetings or seminars. The room is also fitted with Bluetooth sensors to provide additional information for the database about the human researchers present when a particular research data item is presented. The agents, like human users, can update the database with new research items or access existing data already in the database. This setup is illustrated in Figure 3.

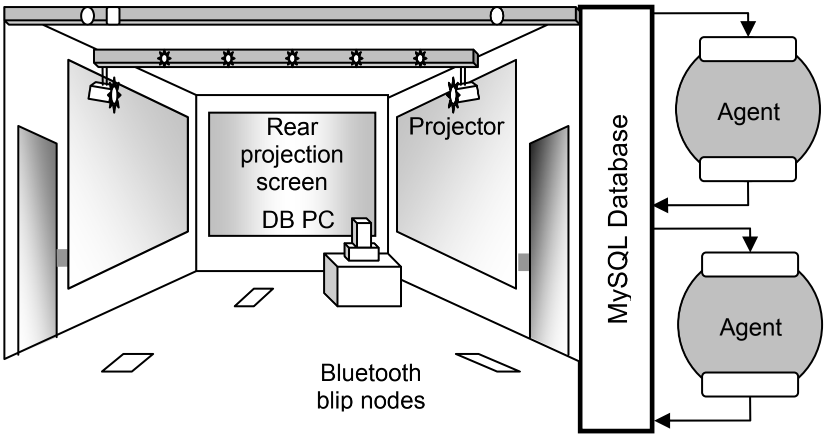

Figure 3. The curious research space agent society and its environment.

Two types of agent models are incorporated into the curious research space: reflex agents and motivated agents. Reflex agents monitor their environment using sensors and reason about the environment using two characteristic processes: sensation and activation as shown in Figure 4(a). The sensation process transforms raw sensor data into structures for further reasoning. The activation process uses a set of rules to choose actions based on sensed data. Rules in reflex agents may be implemented in a number of different ways including as a sequence of programmed operations or as a set of if-then rules forming a grammar. We use both approaches for reflex agents in the curious research space: web-search agents use sequences of programmed operations while presentation agents use a grammar comprising rules for effective presentations. The agent architecture provides a uniform interface of sensors and effectors by which these different types of agent interact with a common environment.

Motivated agents incorporate a motivation process into a reflex agent framework as shown in Figure 4(b). The purpose of the motivation process is to influence the selection of actions through computational models of motivation such as interest, curiosity and novelty in order to facilitate more adaptive, emergent behaviour. Motivation is computed based on an agent's experiences, so, unlike standard reflex agents that exhibit the same behaviour in the same situations, motivated agents may exhibit different behaviour in the same situation based on their previous experiences in their environment. This is the key to their adaptive behaviour.

Curious research space agents can make use of two types of sensor and one effector. Different agents use different combinations of these sensors and effectors and use different reasoning processes to trigger effectors based on sensed data. Agents in the curious research space can sense two types of data from their environment: information about research data presented by humans and information about human presence in the physical space via Bluetooth information. The research data sensor senses URLs in the database, uses PDFbox (http://www.pdfbox.org/) to extract text from the referenced document and creates a word vector (Raghavan and Wong 1986). Word vectors contain a list of the words in a document, excluding stop-words like 'the' and 'and', together with a count of the number of occurrences of each word. The word vector, document title, URL and author are passed to the agent for further reasoning.For example, a word vector for the title of this section might look like:

<curious:1><research:1><space:1>

The Bluetooth sensor monitors human presence in the physical meeting room by monitoring the presence of Bluetooth devices. This allows agents to reason not only about research data, but about the people to whom it may be relevant.

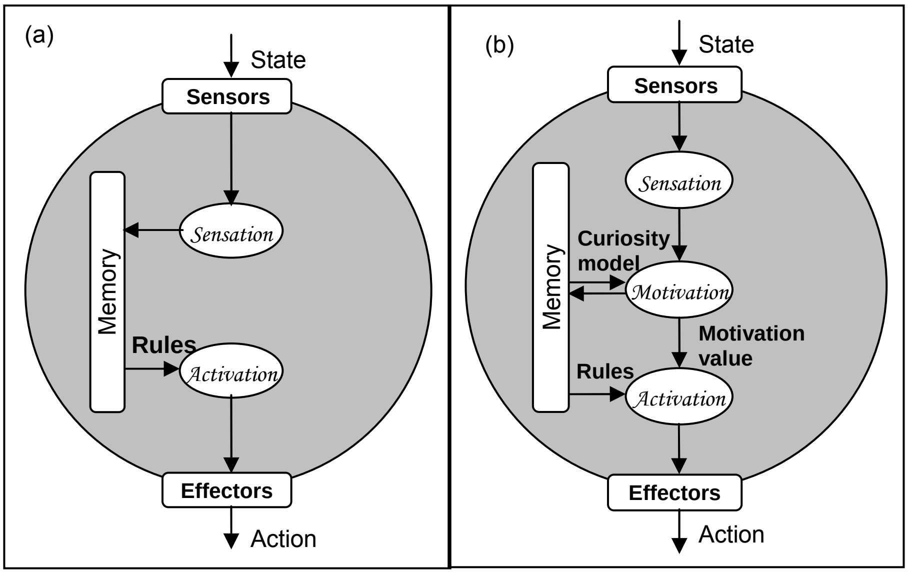

Figure 4. (a) Reflex agents and (b) motivated agents.

4.1. WEB-SEARCH AGENTS

Web search agents are reflex agents that sense human research data and search for new, related data using existing web-search tools. They use research data sensors and research data effectors. The key difference between these agents and existing web-search tools is that key-words used in the search are extracted automatically from human research data associated with the physical environment. We have experimented with two different web-search agents using different approaches to key-word extraction and search. The first web-search agent, a simple search agent, uses the four most frequent words in the feature vector of each human research item sensed to trigger a search with the Yahoo! search engine. The second web-search agent uses the Carrot2 clustering engine to cluster human research data then uses the cluster headings to trigger a search with the Yahoo! search engine. The web-search agents modify the environment by inserting records into the research database. Only search results that are in pdf format are added to the database as an initial means of focusing the search results on files most likely to represent research papers.

4.2. RESEARCH ASSISTANT AGENTS

Research assistant agents are motivated agents that reason about data produced by both human researchers and search agents to identify research data produced by search agents which may be of interest to human researchers. Research assistant agents use research data sensors and Bluetooth sensors to monitor human research data and human presence when research data is added to the research database. They use research data effectors to add records to the database describing research items they computed to be highly novel. We have experimented with a number of different techniques for modelling the motivation component of these agents as a computational model of curiosity. The key idea of all of these techniques, illustrated in
Figure 5, is to cluster or classify human research data using a machine learning or document classification algorithm.

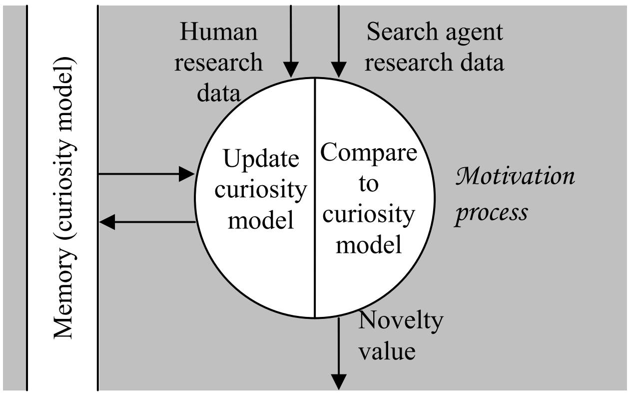

Figure 5. Modelling curiosity as novelty.

Curiosity is modelled as novelty, with the novelty of research data found by search agents computed as the classification error if a search agent document were to be classified, without actually classifying those documents. That is, research data found by search agents does not contribute to the model of research data produced by humans, but is only compared to that model. We experimented with a variety of supervised and unsupervised learning approaches to document classification including K-Means clustering, KNN classification, Naïve-bayes classification and LSI clustering as the motivation component. These include both semantic and word vector based approaches to document classification.

4.3. PRESENTATION AGENTS

Presentation agents are reflex agents that monitor research data produced by research assistant agents and use grammars to produce new research data as presentations of data found by the research assistant agents. Presentation agents also monitor the physical space to determine which presentations to display based on presence information from the environment.

Presentation agents use a research data sensor and a research data effector. The grammars, which define the reasoning process of a presentation agent, describe principles for good presentation design and automate the production of presentations.

At present, we have implemented one presentation agent that produces slides by extracting keywords from interesting documents and laying them out as a single bullet list. We used the s5 slide show system (http://meyerweb.com/eric/tools/s5/) as a template for these slides. In future we hope to extend the techniques used to extract summary data from documents, to include full sentences and images. Likewise, we will continue to develop the layout grammar to extend to these cases.

4.4. EVALUATION OF DIFFERENT AGENT MODELS

We propose to formally evaluate curious research spaces comprising different agents by running lunchtime meetings in which the presentations produced by presentation agents are displayed. However, prior to this, we have some anecdotal evaluations of the research data produced by different agents. The simple search agent, for example, adds many more research data items to the database in response to a batch of human data than does the clustering search agent. However the data found by the simple search agent tends to span a wider range of topics, some of which do not appear to be relevant to the research group being monitored. This is because the simple search agent bases its search on key-words which do not capture the semantics of the document they represent. In contrast, the clustering search agent searches using cluster headings created after a semantic analysis of the human research data.

With respect to the research assistant agents, the K-Means and LSI based approaches to motivation appear to compute interest values which identify as interesting, sets of documents which conform most closely to a human assessment of the same data. KNN clustering does not achieve high enough resolution of interest values, particularly on small sets of human input data. Naïve-bayes as an approach to modelling motivation proved inappropriate as the confidence value produced by classification tends to be universally close to 1 so novelty is high for all documents.

Other approaches to modelling curiosity may also be possible. For example, curiosity can be modelled as interest by applying the Wundt curve to novelty values. In this model, interesting documents are those with a moderate degree of novelty. In this case, however, the question is raised as the values of novelty at which the Wundt curve should peak.

5. Conclusions

The development of motivated and learning agents as an infrastructure for proactive design elements promises to change the way in which we conceive of and design buildings and public spaces. Using sensors to observe the activities in the room and effectors to make changes in the visible aspects of the room opens the possibility of a room as a proactive environment that exhibits characteristics of curiosity. The applications described in this paper highlight some of the alternative computational models and the impact of using different models in different scenarios.

Acknowledgements

This research was supported by an Australian Research Council Discovery Grant titled “Curious Places: Agent-Mediated Self-Aware Worlds”. This paper was written while Mary Lou Maher was working at the National Science Foundation in the United States. Any opinion, findings and conclusions or recommendations expressed in this chapter are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

Berlyne, DE: 1960, *Conflict, arousal and curiosity,* McGraw-Hill, New York.

Brooks, RA et al: 1997, The intelligent room project, *The Second International Cognitive Technology Conference* (CT' 97), Aizu, pp. 271-279.

Coen, M: 1998, Design principles for intelligent environments, *The Fifteenth National / Tenth Conference on Artificial Intelligence / Innovative Applications of Artificial Intelligence*, Madison, Wisconsin, USA, pp. 547-554.

Dieterich, H, Malinowski, U, Khme, T and Schneider-Hufschmidt, M: 1993, State of the art in adaptive user interfaces. *in* M Schneider-Hufschmidt, T Khme and U Malinowski (eds), *Adaptive User Interfaces: Principlse and Practice*, North Holland, Amsterdam.

Graesser, A, VanLehn, K, Rose, C, Jordan, P. and Harter, D: 2001, Intelligent tutoring systems with conversational dialogue, *AI Magazine* **22**(4): 39-52.

Hanssens, N, Kulkarni, A, Tuchinda, R. and Horton, T: 2002, Building agent-based intelligent workspaces, *The Third International Conference on Internet Computing*, pp. 675-681.

Johanson, B, Fox, A and Winograd, T: 2002, The interactive workspaces project: experiences with ubiquitous computing rooms, *IEEE Pervasive Computing* **1**(2): 67-74.

Kulkarni, A: 2002, Design principles of a reactive behavioural system for the intelligent room. Bitstream: *The MIT Journal of EECS Student Research*: 1-5.

Merrick, K and Maher, ML: 2006, Motivated reinforcement learning for non-player characters in persistent computer game worlds, *International Conference on Advances in Computer Entertainment Technology*, CA, USA, pp. (electronic proceedings).

Raghavan, VV and Wong, SKM: 1986, A critical analysis of vector space model for information retrieval, *Journal of the American Society for Information Science* **37**(5): 279-87.

Roman, M et al: 2002, A middleware infrastructure for active spaces, *IEEE Pervasive Computing* **1**(4): 74-83.

Rummery, GA and Niranjan, M: 1994, *On-line q-learning using connectionist systems.* CUED/F-INFENG/TR 166, Cambridge University, Engineering Department.

Saunders, R: 2001, *Curious design agents and artificial creativity.* PhD Thesis, University of Sydney, Sydney.

Stanley, JC: 1976, Computer simulation of a model of habituation, *Nature* **261**: 146-148.

Watkins, C and Dayan, P: 1992, Q-learning, *Machine Learning* **8**(3): 279-292.

Wooldridge, M and Jennings, N: 1995, Intelligent agents: theory and practice, *The Knowledge Engineering Review* **10**(2): 115-152.

Wundt, W: 1910, *Principles of physiological psychology*, Macmillan, New York.