

Intrinsically Motivated Intelligent Rooms

Owen Macindoe and Mary Lou Maher

School of Information Technologies
Faculty of Architecture
University of Sydney, Australia
macindoe@usyd.edu.au, mary@it.usyd.edu.au

Abstract. Intelligent rooms are responsive environments in which human activities are monitored and responses are generated to facilitate these activities. Research and development on intelligent rooms currently focuses on the integration of multiple sensor devices with pre-programmed responses to specific triggers. Developments in intelligent agents towards intrinsically motivated learning agents can be integrated with the concept of an intelligent room. The resulting model focuses developments in intelligent rooms on a characteristic reasoning process that uses motivation to guide action and learning. Using a motivated learning agent model as the basis for an intelligent room opens up the possibility of intelligent environments being able to adapt both to people's changing usage patterns and to the addition of new capabilities, via the addition of new sensors and effectors, with relatively little need for reconfiguration by humans.

1 Introduction

Developing intelligent rooms, such as The Sentient in the Key Centre for Design Computing and Cognition at Sydney University pictured in Fig 1, has been dominated by the development of specific triggers and responses that require a great deal of configuration. In their seminal papers on IE design, Brookes [1] and Coen [2] argued that a key design goal for developing IEs is to enable them to adapt to, and be useful for, everyday activities. Since these papers were published there has not been any published attempt to address the question of how to go about developing such a self-adapting IE. In addition, Brookes and Coen's group at MIT found that configuring new sensor and effector systems to allow their IEs to produce useful behaviours was time consuming and labour intensive. If an IE could learn from its patterns of usage and learn for itself how to manipulate its sensors and effectors usefully both of these issues could be addressed.



Fig. 1. The KCDCC's intelligent environment, The Sentient

A motivated learning agent, recently introduced by Singh et al [3], is an agent that is self-motivated to learn. Self-motivated learning provides a good basis for an intelligent room because we want these environments to be driven to adapt to the installation of new sensors and effectors and changing usage patterns. In this paper we bring together ideas from intelligent environments and computational models of motivation to develop a intrinsically motivated intelligent room.

2 Intelligent Environments

An IE is a physical space for living or working in that is agent controlled and can bring computational power embedded within it to bear in a manner that helps users of the environment perform their daily tasks. The term Intelligent Environment has not been universally adopted and IEs also go under other names such as Jeng's [4] Ubiquitous Smart Spaces. An IE would necessarily need to be able to sense what is happening inside of it and respond to it with effectors - whether lights, projectors, or doors - in order to exhibit intelligent behaviour and help users.

IE research could be regarded as a sub-field of ubiquitous computing in as much as a major aim of ubiquitous computing is to seamlessly integrate computers into everyday living. IEs have several specific design requirements. Brooks [1] and Coen [2] have argued that IEs should adapt to, and be useful for, ordinary everyday activities; they should assist the user, rather than requiring the user to attend to them; they should have a high degree of interactivity; and they should be able to understand the context in which people are trying to use them and behave appropriately. An IE is essentially, as Kulkarni [5] suggests, an immobile robot, but its design requirements differ from those of normal robots, in that it ought to be oriented towards maintaining its internal space rather than exploring or manipulating its environment.

MIT's intelligent room prototype e21, shown in Fig 2, facilitates activities via a system called ReBa, described by Hassens et. al. [7] which is the context handling component of the room. ReBa 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 which responds to a user action and performs a reaction, such as turning on the lights when a user enters the room. Behaviours can then layer on top of one another based on the order of user actions, acknowledging differences in context such as showing a presentation in a lecture setting versus a showing one in an informal meeting. Although ReBa can infer context in this way, it cannot adapt to new ways of using the room. In order for an entirely new context to be created, ReBa's behaviour agents must be pre-programmed to recognize the actions of the user and take an appropriate action. It does not self-adapt to new usage patterns. Furthermore, when new sensors are added to the room, the existing rules must be modified manually if they are to take advantage of the new sensor information. Our model, by contrast, uses a domain in-

dependent mechanism to learn behaviours rather than having the behaviours implemented as part of the agent



Fig. 2. MIT's Intelligent Room Prototype e21, from [6]

Other researchers have taken approaches to designing environments that are not explicitly agent-based. Both the University of Illinois' Gaia [8] and Stanford University's Interactive Workspace Project [9] have taken a more OS-based approach, developing Active Spaces and Interactive Workspaces respectively, which focus on the role of the room as a platform for running applications and de-emphasizing the role of the room as a pro-active facilitator. The specification of an action in these systems is triggered by the user and the behaviour is programmed as an applications developer. Gaia's context service provides the tools for applications developers to create agent-based facilitating applications, and the overall model is reactive rather than adaptive. Georgia Tech's Aware Home Research Initiative plans on incorporating an infrastructure for developing context-aware applications [10], but so far no systems exist which allow IEs to self-adapt to new usage patterns. We believe that a motivated agent-based approach is necessary to allow for this kind of adaptivity.

3 Motivated Learning Agents

In AI literature an agent is anything that can be viewed as perceiving its environment through sensors and then acting within its environment using effectors on the basis of this sensor input. Agent models have a lot in common with IEs: 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 over time with differing ways of mapping sensor input to effector output, from simple rule-based reactive agents through to complex cognitive agents that try to maintain, and reason about, an internal model of the world. The question then is, what kind of agent model would be a suitable basis for an IE?

An IE needs to be driven to assist users, adapt to changes in its configuration, adapt to changing uses of the IE, and understand context. Drives of this kind have been

modelled by the concept of motivation in agent research, leading to several different varieties of motivated agent models. Norman and Long [11, 12, 13] developed a motivated agent model where motivation was modelled by the temporal urgency of tasks to be completed in order for a motivated agent-controlled warehouse to fill orders. Part of the model is shown in Fig 3, which illustrates how the motivation component directed the reasoning process to create new goals for the agent.




Fig. 3. Norman and Long’s motivation model, from [].

Beaudoin and Sloman [14] developed a simulation of a robot nursery in which a robot nursemaid implementing a motivated agent model was shown to effectively prioritise tasks using a sophisticated model of motivation that included logical propositions, temporal urgency, and levels of insistence. In their design of an agent-controlled water filtration plant, Aylett et al. [15] explicitly extended the role of motivation in their agent model to planning, which showed promise despite the relatively simplistic motivation model used. Kasmarik et al [16] experimented with a domain independent model of motivation and used it as a trigger for reinforcement learning in different domain applications.

The requirement for adaptivity in an IE suggests drives towards learning how to behave and how to utilize sensors and effectors, which are themselves instances of the more general drive of curiosity. Saunders and Gero [17] modelled curiosity computationally as a process that internally generates reinforcement signals that reward the discovery of novelty. They then modelled novelty as the property of being similar enough to other entities of the same class so as to be recognisable as part of that class, but different enough from the norm of that class’ form to be unusual. Computationally, novelty was modelled using a self-organizing map that categorized entities presented to the curious agent. The further from the centroid of a class that the new entity’s properties lay in the map, the more novel it was considered, but if it were more than a certain threshold away from the centroid the degree of novelty fell off following a Wundt curve, shown in Fig 4, representing dissatisfaction with an entity’s “strangeness”.




Fig. 4.The Wundt curve. Motivation rises and then falls off as novelty increases. From [16].

They then set out to investigate the utility of this model by using it to simulate the formation of cliques in artistic communities [18], to explore the design space of a simple architectural problem [19], and to provide a richer social force model of human crowds in museums [17]. Their agent models lacked a deep reasoning process but were nevertheless motivated agents, and this model of curiosity as a motivation could be applied to more complicated agent models, such as that which would be required to provide an IE's functionality.

Curiosity by itself is not enough to bring about adaptivity; learning must be a component too. Schmidhuber [18] and Singh et al [3] have developed agent prototypes that are motivated by their own models of curiosity. Schmidhuber developed an agent with a co-evolutionary learning strategy using a highly idiosyncratic model of curiosity that showed promising empirical results in performing exploration when compared with other learning strategies. Singh et. al. developed a model of an intrinsically motivated reinforcement-learning agent illustrated in Fig 5. Inspired by a psychological definition of intrinsic motivation, which is being motivated to do something because it is inherently enjoyable, they developed a learning algorithm in which the learner is rewarded internally for discovering new properties of its domain. They also gave the agent the capacity to build incrementally upon the list of actions that it discovered it could undertake in its domain and allowed them to be chained together into more complicated actions. A comparison between their prototype and a regular reinforcement learning agent showed that it was significantly faster at learning new behaviours. The most interesting feature of the prototype that they built was that the agent was able to learn new behaviours relatively quickly with no human intervention at all. The successes of Schmidhuber and Singh et al.'s motivated learning agents suggest that a motivated learning agent model could be a viable solution to providing the adaptivity required for an IE.

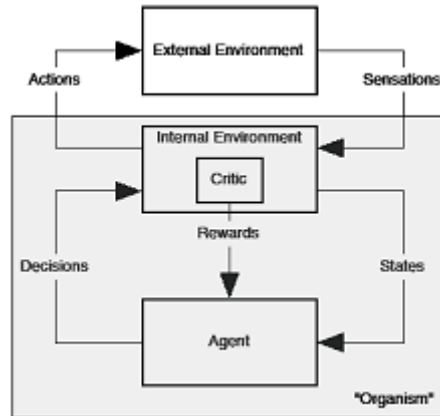


Fig. 5. Singh et. al.'s intrinsically motivated reinforcement learning model. From [3].

4 Intrinsically Motivated Intelligent Room

Combining the ideas in IEs with motivated learning agents leads to a model for an intrinsically motivated intelligent room. Motivation can play a valuable role in the agent model for an intelligent room generally, not just in learning, because it provides a model for the pro-active characteristics that are desirable in IEs. We present a motivated agent model for an intelligent room that combines a self organising map for characterising the clusters of events in the room, with the characteristic Wundt curve being used to identify interesting novel events, and data mining capabilities to find patterns of behaviour in the states of the room that correspond to a selected interesting event. This model is shown in Fig 6. The model assumes two significant entities: the world and the agent. The world is described at any point in time by the data that can be sensed in the intelligent room. The agent has sensors to sense the state of the world, effectors to change certain aspects of the state of the room, a memory of world states and events, and a reasoning process that includes motivation, action, and learning.

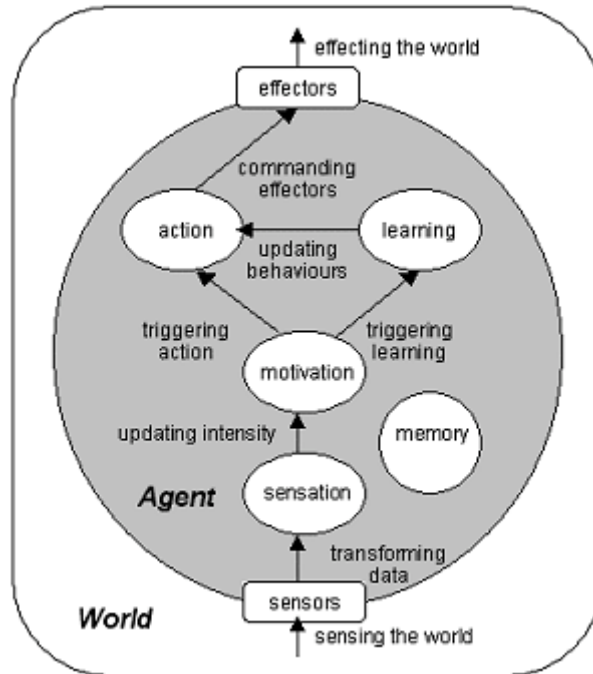


Fig. 6. The intrinsically motivated intelligent room model.

4.1 The World State

The motivated learning agent exists within a specific world. The state of the world is the basis for agent's interaction with the world; therefore it becomes the basis for configuring sensors and effectors and adapting to new behaviour patterns. The world state at time t , $W(t)$, is characterised as a partitioned tuple of sensor inputs, which are in turn represented as attribute-value pairs such as `PRESSURE_PAD=ON`. One side of the partition represents inputs from sensors without associated effectors, such as a pressure pad in the floor. The other side of the partition represents inputs from sensors that do have associated effectors, such as a sensor attached to a light switch which can be both activated manually by a human operator and automatically by the room itself. A world state representation $W(t)$ will therefore take the form:

```
W(t) ::= <senseData>
<senseData> ::= "(" <senseOnly> "|" <senseEffect> ")"
```

And an example of such a representation is:

```
W(0) = (PRESSURE_PAD=ON | LIGHT_DIMMER_INTENSITY=0.5,
DESK_LAMP=ON)
```

This distinction is relevant because the intention is for the motivated agent to learn behavioural rules via data mining, so any relevant rule must include changes in the ef-

fectable sensor data part of its sufficient conditions. For instance a rule such as the following would represent a behaviour that the IE would not have the capacity to enact since it does not have the effectors necessary to achieve it:

```
IF SENSE = (PRESSURE_PAD=ON) THEN EFFECT =  
(PRESSURE_PAD_4=ON)
```

4.2 Sensation

In the sensation process, sensor input from the world is converted into a form suitable for performing reasoning and learning. The new world state $W(t)$ is stored in the set S of all world states sensed by the agent. The sensation component also records events or changes in the world state. An event is represented as $\Delta(t)$, the changes in sensor inputs between $W(t)$ and $W(t-1)$. $\Delta(t)$ takes the same form as $W(t)$, a partitioned tuple, but the values of the tuple represent the change in value between $W(t)$ and $W(t-1)$ with numeric values being calculated as normalized differences and nominal elements being 0 if no change occurred and 1 if one did occur. For example:

```
W(0) = (PRESSURE_PAD=ON | LIGHT_DIMMER_INTENSITY=0.5,  
        DESK_LAMP=ON)  
  
W(1) = (PRESSURE_PAD=OFF | LIGHT_DIMMER_INTENSITY=0.8,  
        DESK_LAMP=ON)  
  
 $\Delta(1)$  = (PRESSURE_PAD=1 | LIGHT_DIMMER_INTENSITY=0.3,  
           DESK_LAMP=0)
```

The sensation component is also responsible for handling sensor faults in order to protect the integrity of the world representations so that reasoning and learning processes are not adversely effected by hardware faults.

4.3 Motivation

The intelligent room is expected to respond to human actions within reasonable time bounds, therefore it is desirable to only perform learning when it is useful to do so. Since the problem of self-adaptivity centers upon adapting to novel usage patterns and novel sensor and effector inputs from newly installed sensors and effectors, “useful” times to learn here correspond to times when novel world states occur. Responding to novel world states is governed in this model by a motivational model of curiosity.

The motivation component of the agent model directs attention to action or learning as appropriate. The motivation component is a model of curiosity similar to the Saunders and Gero model outlined in section 3, which maps $\Delta(t)$ to a motivational intensity, $I(t)$. This is done via a comparison between $\Delta(t)$ and previous $\Delta(t')$ s, where $t' < t$, using a self organizing map and Wundt curve as in the Saunders and Gero model. If $I(t)$ is greater than a required threshold, then the curiosity’s intensity triggers the learning component of the agent model, otherwise the action component of the agent is triggered.

4.4 Learning

Because it is inappropriate for an intelligent room to experiment with changes in the state of the room, learning must rely upon drawing inferences from previously experienced world states without being able to effect the environment during the learning process. The aim of the learning component of the agent model is to infer a set R of behavioural rules from the set of stored world data S and then store R in memory for the action component to utilize. Such behavioural rules will be of the form:

Rule ::= IF SENSE = <window> THEN EFFECT = <action>

Where <window> is a tuple of <sensor-data> and <action> is a tuple of attribute-value pairs consisting only of effectable sensor data. Such rules are formed by considering the changes in world state within a given time window and constructing rules to enact equivalent changes when sufficient support and confidence levels exist for such a rule to be derived. Data mining techniques such as Generalised Sequential Pattern (GSP) mining can find these rules from the memory of world state tuples.

4.5 Action

The action component of the agent model maps the most recently sensed world state $W(t)$ and previous world states within a given time window to a rule from the set of behavioural rules R to be executed by the IE's effectors. It then sends the appropriate commands to the IE's effectors to enact the changes in the world dictated by the rule selected.

4.6 Memory

The sensation, motivation, learning, and action components all require information about earlier states of the world, and all bar action update that information. This necessitates a memory component for storing previous worlds states, deltas, and behavioural rules. The various kinds of interactions with memory are outlined in the sections on the specific components and Fig 7 provides a diagrammatic summary.

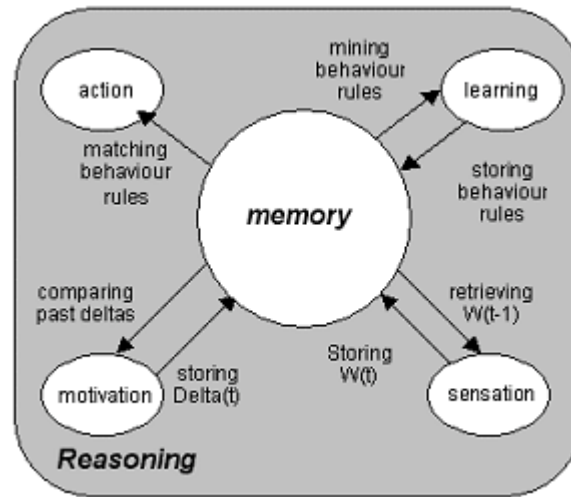


Fig. 7. Interactions between the reasoning and memory components of the model.

5 Conclusions

A model for an intelligent room based on an intrinsically motivated learning agent moves us closer to an adaptable intelligent environment. Our initial tests with this model include sensor data that identifies different behaviours associated with the location of people in the room (the pressure pads) and the state of the electric devices in the room (lights, projectors, applications being projected). Given this kind of data, behaviours can be learned that are based on patterns of use, rather than on the identities of the individuals in the room. The validation of this model is a test of its adaptability, that is, can the room change its behaviour when new sensors or effectors are added without requiring additional programming.

References

1. Brooks, R.A., Coen, M., Dang, D., DeBonet, J., Kramer, J., Lozano-Perez, T., Mellor, J., Pook, P., Stauffer, C., Stein, L., Torrance, M., Wessler, M.: The Intelligent Room Project. In: Proceedings of the Second International Cognitive Technology Conference (CT'97). Aizu, Japan (1997) 271-279
2. Coen, M.H.: Design Principles for Intelligent Environments. In: Proceedings of the Fifteenth National / Tenth Conference on Artificial Intelligence / Innovative Applications of Artificial Intelligence. Madison, Wisconsin, United States (1998) 547-554
3. Singh, S., Barto, A.G., and Chentanez, N.: Intrinsically Motivated Reinforcement Learning. <http://www.eecs.umich.edu/~baveja/Papers/FinalNIPSIMRL.pdf> Accessed 7/4/2004 (2004)
4. Jeng, T.: Designing a Ubiquitous Smart Space of the Future: The Principle of Mapping. In: Gero, J.S. (ed.): Design Computing and Cognition '04. Kluwer Academic Publishers, The Netherlands (2004) 579-592

5. Kulkarni, A.: Design Principles of a Reactive Behavioral System for The Intelligent Room. In: Bitstream: The MIT Journal of EECS Student Research, April 2002 Edition. Cambridge, MA (2002) 1-5
6. Kottahachchi, B., Laddaga, R.: Access Controls for Intelligent Environments. In: Proceedings of ISDA '04: 4th Annual International Conference on Intelligent Systems Design and Applications. Budapest, Hungary (2004)
7. Hanssens, N., Kulkarni, A., Tuchinda, R., Horton, T.: Building Agent-Based Intelligent Workspaces. In: Proceedings of the 3rd International Conference on Internet Computing, (2002) 675-681
8. Román, M., Hess, C.K., Cerqueira, R., Ranganathan, A., Campbell, R.H., Nahrstedt, K.N.: Gaia: A Middleware Infrastructure to Enable Active Spaces. In: IEEE Pervasive Computing Oct-Dec (2002) 74-83
9. Johanson, B., Fox, A., Winograd, T.: The Interactive Workspaces Project: Experiences With Ubiquitous Computing Rooms. In: IEEE Pervasive Computing Magazine, vol. 1, no. 2, April-June (2002) 67-74
10. Kidd, C., Orr, R., Abowd, G., Atkeson, C., Essa, I., MacIntyre, B., Mynatt, E., Starner, T., Newstetter, W.: The Aware Home: A Living Laboratory for Ubiquitous Computing Research. In: Proceedings of the Second International Workshop on Cooperative Buildings. Pittsburgh, PA (1999)
11. Norman, T.J., Long, D.P.: Goal Creation in Motivated Agents. In: Wooldridge, M.J., Jennings, N.R. (eds.): Intelligent Agents: Proceedings of the First International Workshop on Agent Theories, Architectures and Languages, Volume 890 of Lecture Notes in Artificial Intelligence, Springer-Verlag (1995a) 277-290
12. Norman, T.J. and Long, D.P.: Alarms: Heuristics for the Control of Reasoning Attention. In: Moore, J. and Lehman, J. (eds.): Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society, Pittsburgh, PA (1995b) 494-499
13. Norman, T.J. and Long, D.P.: Alarms: An Implementation of Motivated Agency. In: Wooldridge, M.J., Müller, J.P., Tambe, M. (eds.): Intelligent Agents II: Proceedings of the Second International Workshop on Agent Theories, Architectures, and Languages, Volume 1037 of Lecture Notes in Artificial Intelligence, Springer-Verlag, (1996) 219-234
14. Beaudoin, L.P., Sloman, A.: A Study of Motive Processing and Attention. In: Sloman, A., Hogg, D., Humphreys, G., Partridge, D., Ramsay, A. (eds.): Prospects for Artificial Intelligence. IOS Press, Amsterdam (1993) 229-238
15. Aylett, R.S., Coddington, A.M., Petley, G.J.: Agent-Based Continuous Planning. <http://planning.cis.strath.ac.uk/plansig/pastsigs/open-univ-19/aylett.pdf> Accessed 7/4/2005 (2000)
16. Kasmarik, K., Uther, W., Maher, M.L.: Motivated Agents. Poster in: IJCAI (2005)
17. Saunders, R., Gero, J.S.: Designing for Interest and Novelty: Motivating Design Agents. In: de Vries, B., van Leeuwen, J., Achten, H. (eds.): CAADFutures 2001. Kluwer, Dordrecht (2001) 725-738
18. Saunders, R., Gero, J.S.: Curious Agents and Situated Design Evaluations. In: Gero, J.S., Brazier, F. (eds.): Agents in Design 2002. Key Centre of Design Computing and Cognition, University of Sydney, Australia (2002) 133-149
19. Saunders, R., Gero, J.S.: The Digital Clockwork Muse: A Computational Model of Aesthetic Evolution. In: Wiggins, G. (ed.): Proceedings of the AISB'01 Symposium on AI and Creativity in Arts and Science. University of York, York, UK (2001) 12-21
20. Schmidhuber, J.: What's interesting? Technical Report IDSIA-35-97. (1997) Lugano, Switzerland