CAN DESIGNS THEMSELVES BE CREATIVE?[1]

MARY LOU MAHER
Key Centre of Design Computing and Cognition, Australia
National ICT Australia

KATHRYN MERRICK
University of Sydney, Australia
National ICT Australia

OWEN MACINDOE
University of Sydney, Australia

Abstract. Studies of design and creativity usually investigate the processes that result in a design product that is distinguished as being a creative product. The focus of these studies comprises the definition and recognition of creativity and the development of cognitive and computational models of creative design. We consider these studies in light of the development of a design that is itself creative. If some of the characteristics associated with creative design or a creative product can be used to describe the behaviour of a product, then we assert that the design is itself creative. We embed a computational model of curiosity within the design of specific product in order to ascribe the characteristics of creativity to the product. We illustrate this concept with the design of a curious place.

1. Introduction

We introduce the concept of a product that has an embedded computational process capable of being creative. Our implementation of these ideas is grounded in the concept of an intelligent room. In our intelligent room, the room learns new behaviours by looking for novel events and creating goals to learn more about how to predict and achieve those events with the purpose of aiding the human activities within the room. We propose that the resulting

---

[1] Published in the Proceedings of The Sixth International Roundtable Conference on Computational on Cognitive Models of Creative Design, http://www.arch.usyd.edu.au/kcdc/conferences/hi05

room is creative, in the sense that it is able to detect novelty and to adapt its behaviour in response to a novel event.

When we consider computational or cognitive models of creative design, we study or develop processes and mechanisms that lead to creative designs. This assumes a common understanding of the difference between creative design and routine design as well as a distinction between a creative process and a creative product.

It is important in a discussion on computational models of creative design to distinguish creative design from other kinds of design. This is necessary or all models of design are models of creative design. The qualifier "creative" makes reference to something novel or unexpected. The idea of novelty is relative. A specific product may be familiar to one person, and therefore not considered creative, while to another person the same product may be novel and considered creative. Boden (1994) talks about two types of creativity: one is psychological (P-creativity), "that is the creative idea is apparent to the person in whose mind it arises", and the other is historical (H-creativity), "that is the creative idea is P-creative and no one else, in all human history, has had it before". She states more strongly that "many creative ideas [are interesting] as they concern novel ideas that not only *did not* happen before, ... but that *could not* have happened before".

Dasgupta (1994) expands this notion in terms of psychological novel (PN-creative), psychological original (PO-creative), historical novel (HN-creative) and historical original (HO-creative). His Computational Theory of Scientific Creativity (CTSC) attempts to encompass agents A and cognitive processes P for which PO-creativity and HO-creativity is possible respectively for the domain and the community. This *explains* rather than predicts P conducted by A in terms of a knowledge level process P(KL) to transform a scientific goal into a solution. This distinction between novel and original, or personal and historical, is important when developing a model for a creative design. In our model, we focus on personal creativity and novelty, where the design is embedded with an agency that acts and learns to respond to novel events with no consideration for a social construct of novelty or evaluation.

Gero and Maher (1993) distinguish between routine, innovative, and novel design using the concept of design variables to make the distinction. In a routine design, the variables and the values associated with the design are known in advance. In innovative design, the variables are known but some of the values for the variables fall outside the known range. In novel design, the designer introduces new variables, defining a new kind of design that was not part of the original search space. In our model, we have a reasoning process in which the agent generates it own goals and behaviours, and therefore appeals to Gero and Maher's definition of novel design.

Finally, the distinction between a creative process and a creative product indicates what is being evaluated. A creative process is one in which the process that generates the design is novel, and a creative product is when we evaluate the result of the design process as being novel. In our model, we consider a process as part of a product. The product has agency, and that agency produces creative behaviours. Therefore, the product is creative because it has creative behaviours.

There are numerous other characteristics associated with creative design beyond novelty. These include aesthetic appeal, quality, unexpectedness, uncommonness, peer-recognition, influence, intelligence, learning, and popularity (Runco and Pritzker 1999). This more broadly defines creative design as being more than just novel, and including judgements related to its appeal and usefulness. In our model we focus on novelty, and although we don't preclude the other characteristics in an evaluation of our product, we do not incorporate them in the computational model.

In this paper, we consider a computational model of curiosity as a basis for developing a product whose behaviour is creative. Our approach starts with the basic concept of an agent model, extended to include motivation driven by curiosity to learn new behaviours. In this paper we include a review of motivation theories as a basis to inform a computational model of motivation and provoke discussion related to models of creativity. We associate a motivated agent with a design product, giving the product behaviours that include curiosity. We propose that the integration of curiosity with agent models as a reasoning component of a design results in a design that responds and adapts to its use.

2. Adaptable Designs by Incorporating Agents into the Design Product

We have been designing, implementing, and inhabiting virtual worlds for education and collaborative design activities using various agent models to give aspects of the world agency (Smith et al 2004; Gu and Maher 2004; Clark and Maher 2005 and Rosenman et al 2005). These 3D networked virtual environments provide a sense of place in which each object in the place comprises a 3D model with location in the 3D world and a computational model that describes the behaviour of the object. Such environments provide a good platform for studying computational models of design because they are a closed world in terms of the objects in the world being knowable by querying the virtual world server, and are open worlds because people can inhabit and interact with the world in unpredictable ways similar to the physical world. By associating objects in the 3D virtual world with an agent model, we can study computational models of design in which the design itself is able to design and adapt to its use.

There is no universal definition for the term agent. However in the context of computer science, agents as intentional systems operate independently and rationally, seeking to achieve goals by interacting with their environment (Wooldridge and Jennings 1995). An agent has the ability to operate usefully by itself, however the increasing interconnection and networking of computers is making this situation rare. Typically, the agent interacts with other agents (Huhns and Stephens 1999). Hence the concept of multi-agent system is introduced with the applications of distributed artificial intelligence.

In object-oriented systems, objects are defined as computational entities that encapsulate some states, are able to perform actions, or methods on this state, and communicate by message passing. There are similarities between agents and objects, but there are also significant differences (Wooldridge 1999). Agents embody a stronger notion of autonomy than objects, and in particular, they decide for themselves whether or not to perform an action on request from another agent. In addition, agents are capable of flexible (reflexive, reactive, reflective/proactive and social) behaviours, and the standard object model has nothing to say about such types of behaviours.

In Figure 1 we show what these virtual worlds look like, and what we mean by an agent model. The left side image shows the 3D virtual world as a place that is inhabited by people represented as avatars. The 3D world is made up of individual objects, each associated with a 3D model and a computational model. The right side of the figure shows an agent model. Each agent has sensors that are able to sense information about its environment, in this case the 3D world, and effectors that are able to change their environment. This combination of virtual worlds and agent models allows us to develop and study designs in which the components of the design have agency. Once we can associate agents with the objects in the design, we can consider whether a design can be creative itself.

Figure 1. Virtual worlds and agents.

Maher and Gero (2002) developed a multi-agent system as the core of a 3D multi-user virtual world. Each object in the world is an agent in a multi-agent system. The agent model provides a common vocabulary for describing, representing, and implementing agent knowledge and communication. The agent can sense its own environment and can generate or modify the spatial infrastructure needed for a specific activity for the users of the world. The common agent model is illustrated in Figure 2, where each agent has five kinds of reasoning: sensation, perception, conception, hypothesizer, and action.

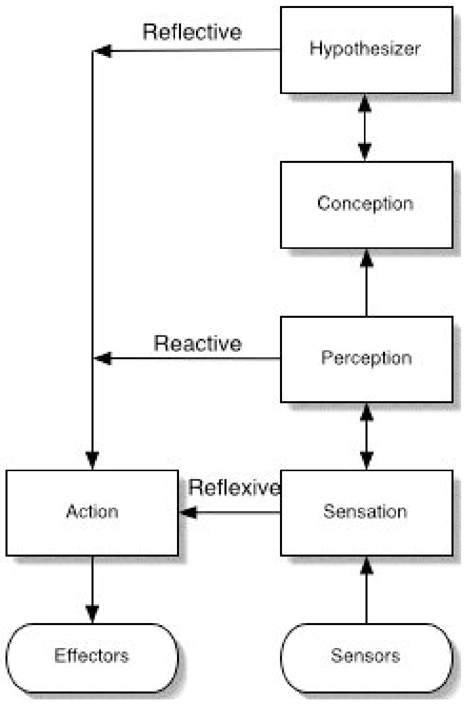

Figure 2. Agent model showing levels of reasoning (Maher and Gero, 2002).

Sensors receive information about the state of the world at any time t. This includes the objects and their properties including location, the avatars and their properties, and itself as an object in the world. **Sensation** transforms raw input from the Sensors into structures more appropriate for reasoning and learning. **Perception** transforms sense-data into the percepts, or perceptual objects, that are used both to interpret interactions and as the units with which concepts are constructed. Percepts are grounded patterns of invariance over interactive experiences, and are constructed by clustering like patterns into equivalence classes so as to partition the sensory representation space. Perception is driven both by concepts and by the sense-data. **Conception** learns and uses concepts about the world to reinforce or modify the agent's beliefs and goals. Concepts are abstractions of experience that confer a predictive ability for new situations. The concept of a meeting, for example, is a representation of the activities of the agent with which meetings are involved, and its various realizations are predictions of possible interaction. **Hypothesizer** identifies mismatches between the current and desired states of the world, and reasons about which goal should be achieved in order to reduce or eliminate that mismatch. It identifies possible actions which when executed will change the world to meet those goals. **Action** reasons about which sequence of operations on the world, when executed, can achieve a specific goal. **Effectors** are the means by which actions are achieved by making changes to the 3D virtual world.

This model was extended by Gu and Maher (2004) to develop a Generative Design Agent (GDA). The GDA responds to its environment by designing places for the avatars needs. In addition to the computational processes in the Maher and Gero model, the GDA includes a design process that is essentially a design grammar. While this grammar is able to add new objects to the world and remove them as determined by a set of intended functions for the virtual world, the process does not specifically address novelty or computational models of creativity.

The agent approach to virtual worlds provides for new kinds of interaction among the elements of the virtual world and between people and the virtual world that makes both the virtual environment and interactions with it dynamic.

3. Motivated Agents

The reasoning process used in the design of adaptable virtual worlds is based on a cognitive agent model. The model describes the levels of reasoning and provides a framework for the knowledge the agent needs in order to respond to changes in the virtual world. However, the agent model assumes that the agent has all of the knowledge needed to respond to the world when the agent is designed and implemented. An alternative is to define an agent model that includes the ability to change its knowledge over time in response to its observation and interaction in the world. Various theories of motivation and motivated learning agents inform a more general computational model of motivated agents.

3.1. MOTIVATION THEORIES

Motivation is the cause of action (Mook 1987). When we ask the question: "Why did he or she do that?" we are inquiring about an individual's motivation. Motivation is though to have three primary functions: a directing function that steers an individual's behaviour towards or away from specific goals, an activating function that energises action in pursuit of goals and an organising function that influences the combination of behavioural components into coherent, goal-oriented behavioural sequences (Green et al 1984; Kandel et al 1995).

Psychological study of motivation searches for theories that describe the functions of motivation in natural systems such as humans and animals. Motivation theorists do not agree on a unified causal explanation of the behaviour of natural systems. Rather, causation has been attributed to such factors as the environment, physiology, the psyche or social forces with early researchers tending to focus on one or the other of these views. Various attempts have been made to either classify (Green et al 1984; Mook 1987) or synthesise (Alderfer 1972; Maslow 1954) the large body of research related to psychological motivation.

While some computational models of motivation have been closely informed by psychological research, others introduce new ideas specifically tailored to artificial systems. In an effort to classify motivation theories in a manner that is relevant for both natural and artificial systems and which does not encroach on the terms used by psychological motivation theorists, we have used three broad categories: biological motivation theories, motivation theories of the mind and social motivation theories.

3.2. MOTIVATION THEORIES FOR NATURAL SYSTEMS

3.2.1. Biological Motivation Theories

Biological motivation theories attempt to explain the motivation processes that work within the biological system of a behaving organism. The **drive theory** of motivation holds that homeostatic requirements drive an individual to restore some optimal biological condition when stimulus input is not congruous with that condition. For example, high body temperature might drive an individual to sweat in order to restore its optimal body temperature. Drive theory was developed in its most elaborate and systematic way by Hull (1943; 1952) as a central part of his theory of behaviour. Hull's theory postulates that behaviour is a response both to a motivational factor called drive and to habits that are reinforced during an individual's lifetime. The more often a response is reinforced, the more habitual it becomes in a particular situation and the more likely it is to be repeated when the conditions are the same as those in which it was reinforced. When a response has become a habit through frequent reinforcement it comes to be performed more intensely under conditions of high drive. Hull modelled the relationship between habit, drive and a behavioural response as multiplicative based on the assumption that both habit and drive are necessary for behaviour.

In recent years there has been a tendency to drop the concept of drive, in particular because it is no longer believed that drives can be considered as unitary variables (McFarland 1995). The notion of high or low hunger, for example, is a misnomer. Hunger is more accurately described in terms of a number of variables such as fat, protein and carbohydrates. As an alternative to drive theory, McFarland proposed the idea of a **motivational state**. An individual's motivational state can be represented by a point in a motivational space. The axes of the space are important motivational stimuli such as fat or protein levels or the strength of some external stimulus. The major difference between the state-space approach to motivation and the drive concept is that the state-space approach makes no assumption that different motivational factors are multiplicative or about the relationship between motivation and behaviour.

In both drive and motivational state theories, a total absence of homeostatic drives such as hunger or thirst should produce an individual that does, and seeks to do, nothing. However, studies of sensory deprivation and isolation in the early 1950s showed that low levels of stimulation, which should produce little drive and hence be attractive according to drive theory, are in fact unattractive and produce a tendency to seek out stimulus complexity. **Arousal theory** offers an alternative to drive theory's explanation of the intensive aspects of behaviour by stating that we seek from our environment, not a universally minimal amount of stimulation, but rather a moderate or optimal level of stimulation so there is an inverted U-shaped relationship between the intensity of a stimulus and its pleasantness (Wundt 1910). A general model for the conditions that produce arousal assumes that arousal is a response to a change in the level of stimulation to which a person is exposed between an existing condition of stimulation and a new and different condition (Green et al 1984). Implicit in the idea of an existing state is the assumption that individuals tend to establish a baseline level of stimulation through constant adjustments and adaptations to their environment. Over time a loss of sensitivity accrues with prolonged exposure to a particular stimulus.

3.2.2 Motivation Theories of the Mind

The view of hunger and thirst as homeostatic drives or optimal arousal theory imply that feeding, drinking or exploration are initiated as a result of monitored changes in physiological state. However, in addition to occurring in response to physiological changes, behaviours such as feeding and drinking often occur in anticipation of such changes (McFarland 1995). In recognition of this, cognitive motivation theories focus on questions such as what determines the consequences of behaviour, how do consequences influence behaviour and to what extent do individuals account for the probable consequences of future behaviour in terms of the costs and benefits of different courses of actions.

The **operant theory** of motivation recognises that individuals are not only driven by deprivations and needs but may be guided to important goals by perceptions and cognitions. When an individual does something that is rewarded, for example, they are not influenced by any real or imagined loss of drive but by the idea of being rewarded. A voluntary response emitted by an individual in order to achieve some reward is called an operant. Skinner's law of conditioning for operants (1938) states that if the occurrence of an operant is followed by the presentation of a reinforcing stimulus, the strength of the operant will be increased. The converse law, the law of extinction, states that if the occurrence of an operant already strengthened through conditioning is not followed by the reinforcing stimulus, the strength is decreased. A reinforcing stimulus, or reward, is not assigned any specific properties other than that it follows an operant. For example, it may be a pleasant, internal feeling of satisfaction, the receipt of money from an external source, an unpleasant, internal feeling of boredom or an electric shock from an external source.

Successful acquisition of reward triggers the formation of mental associations between acts and the rewards that follow them. This association generates an expectancy that if the act is repeated it will be rewarded again. The expectancy of being rewarded after some responses forms the basis of incentive. The **expectancy-value theory** of incentive defines incentive as a multiple of the expectancy of receiving a reward and the value of that reward to the individual.

The notion of the value of a reward to an individual was approached by Atkinson et al (1966; 1974) in their theory of **achievement oriented motivation**. Atkinson defines the value of a reward to an individual in terms of the tendency of that individual to either approach success or avoid failure. If an individual is motivated by a tendency to approach success, Atkinson proposes that they will evaluate the potential reward of a situation in terms of their probability of success so that success on a difficult task is more valuable than success on a simple one.

Some individuals, rather than being motivated to approach success, are motivated simply to avoid failure. If an individual is motivated by a tendency to avoid failure they will evaluate the potential reward of a situation differently to an individual motivated by a tendency to approach success. In an individual motivated to avoid failure, the higher the probability of achieving a task, the greater the negative incentive associated with failure at that task. Thus, individuals motivated to avoid failure tend to choose either easy tasks at which they are likely to succeed or difficult tasks for which there is a clear reason for failure.

While expectancies and values together determine an individual's orientation toward future behaviour, theories such as the expectancy-value theory of incentive do not explain how expectancies and values are formed. Rather an individual's cognitive representation of the environment and their role in it is simply assumed to exist. **Attribution theory** seeks to provide this explanation. A causal attribution is an inference about why some event has taken place. An attribution may be about one's own behaviour or about another's behaviour. Heider (1958) introduced the idea that people follow specifiable rules in interpreting the causes of behaviour. Attribution theory attempts to specify the processes that are involved when an individual develops an explanation the behaviour of others or themselves. Heider used attribution theory to develop his **naïve analysis of action theory** which describes the cause of behaviour in terms of the average person's commonsense analysis of behaviour. Central to Heider's theory was the idea that, along with professional psychologists, naïve perceivers share the belief that there are two classes of causes: personal forces and environmental forces. Heider further subdivided each of these forces into two categories: ability and trying, task difficulty and luck respectively. He proposed that the relationship between ability and task difficulty is additive. That is, environmental forces could oppose or support the personal force and thus increase or reduce its effectiveness. Further he proposed that the personal force, trying, is made up of two components, intention and exertion, and that successful action depends on the presence of both.

Attribution theory can produce an explanation for some behaviours in terms of personal forces such as physiological drives "because they are hungry" and environmental forces such as extrinsic rewards "to earn money". However there are other behaviours that are inexplicable to the average observer. For example, some people make a pastime of skydiving "for fun" or climb mountains "because they are there". These behaviours involve exploration, seeking novelty, curiosity or meeting a challenge. To this extent they are all intrinsically motivating. White (1959) first argued for the existence of what he called **effectance motivation** or competence motivation. He proposed that individuals are motivated to engage in behaviours that can satisfy the desire to feel self-determining and competent. Effectance motivation is the desire to deal effectively with ones environment. White believes that effectence motivation is always present but is only manifested when other more basic needs are satisfied. Specifically, behaviours such as exploratory play, curiosity and stimulation seeking would be expected to appear only when an individual is otherwise homeostatically balanced. White also proposed that effectance motivation is undifferentiated, meaning that the satisfaction of the effectance motive is not tied specifically to any given behaviour. Rather any behaviour that allows an individual to deal effectively with its surroundings can satisfy the motive.

In their theory of **intrinsic motivation**, Deci and Ryan (1985) extended the theory of competence motivation to define the types of behaviours that will permit an individual to gain a sense of competence and self-determination. They proposed that individuals are involved in an ongoing, cyclical process of seeking out or creating optimally challenging situations and then attempting to conquer those challenges. A feeling of competence emerges from situations in which an individual is optimally challenged. Optimal challenging situations are based on an individual's unique complement of skills. The situations that are most intrinsically motivating are those that contain information relevant to structures already stored and mastered but that are discrepant enough to call forth adaptive modifications. Overly familiar or excessively repetitive tasks and tasks that greatly exceed existing capacities will trigger boredom and distress respectively (Hunt 1975). In other words, individuals will orient themselves towards an activity in some domain of behaviour where they are required to learn or stretch their abilities by a small amount, that is, tasks that are neither too difficult nor too easy.

3.2.3. Social Motivation Theories

Where biological motivation theories and motivation theories of the mind are concerned with the individual, social motivation theories are concerned with what causes individuals to act when they are in contact with one another. One of the earliest theories that offers an explanation for motivation in social situations is Charles Darwin's **theory of evolution** (Darwin 1859). Darwin's idea was that animals have the structural and behavioural characteristics required to survive and breed within their habitats. His theory has three key components. Firstly, animals vary from one another within a species. Secondly, animals pass on their characteristics to their offspring. Thirdly, variation within a species means that some members of the species are better adapted than other to the ecology in which they live. Those better adapted are more likely to have offspring and pass on their structural and behavioural characteristics. Those that are poorly adapted will have fewer offspring so their characteristics will diminish over successive generations. Thus the cause of an individual's behaviour can be thought of as influenced by generations of the individual's ancestors and by the selection pressure by the environment in which the species lives.

Alternative social theories of motivation constrain the influence on an individual's behaviour to that individual's social contemporaries rather than their ancestors or environment. In his **theory of cultural effect**, Mook holds that the culture of the society in which an individual lives affects action in two ways (Mook 1987). Firstly he writes that it determines what skills, thoughts and schemata are cognitively available to an individual in a particular situation. For example an individual from western society lost in a forest may not have the notion of eating ants cognitively available as a means of satiating hunger. Secondly, Mook notes that cultural values affect what selections an individual will make from those that are cognitively available. For example suppose someone informed the lost individual that ants are a good source of protein. The individual might still balk at eating them based on their cultural perception of ants as dirty or ugly.

Similar to the theory of cultural effect but confined to even smaller social groups is the notion of **conformity**. The term conformity refers to behaviour that an individual engages in because of a real or imagined group pressure. It must be different from what the individual might have done were the pressure not exerted. Research has shown that conformity pressures can be powerful and effective motivators in both small and large groups.

3.3. MOTIVATION THEORIES FOR ARTIFICIAL SYSTEMS

Motivation theories for artificial systems have been produced by researchers from the artificial life and artificial intelligence communities either to gain insight into the causes of action in living organisms or to construct new artificial systems that have some computational advantage when solving complex problems. Much of the research concerned with creating artificial systems that mimic the biological, cognitive and social properties of systems found in nature is based on the concept of an agent. While some computational models of motivation have been informed by biological and psychological research, others introduce new ideas specifically tailored to artificial systems.

3.3.1. Motivation Theories Derived from Biological Systems

The term artificial life is used to describe research into human-made systems that possess some of the essential properties of life. As a result, artificial life researchers concerned with motivation theories tend to base their models on biological motivation theories. In fact, motivations, according to artificial life researchers (Avila-Garcia and Canamero 2002; Gershenson 2001), constitute urges to action based on internal bodily needs related to self sufficiency and survival. As a result, the computational models of motivation emerging from the artificial life community tend to implement a homeostatic process to maintain essential physiological variables within certain ranges.

Action selection architectures for autonomous agents make decisions about what behaviours to execute in order to satisfy internal goals and guarantee survival in a given environment and situation. A computational model of the drive theory of motivation is one approach to building an action selection architecture (Canamero 1997; Gershenson 2001). Motivations in computational models of drive theory are characterised by a set of controlled essential physiological variables, a set of drives to increase or decrease the level of the various controlled variables, a set of external incentive stimuli that can increase a motivation's intensity and a behavioural tendency of approach or avoidance towards these stimuli. A feedback detector generates an error signal, the drive, when the value of a physiological variable departs from its set-point. This triggers the execution of inhibitory and excitatory behaviours to adjust the variable in the appropriate direction. Each motivation receiving an error signal from its feedback detector receives an intensity or activation level proportional to the magnitude of the error.

3.3.2. Motivation Theories Derived from Theories of the Mind

Artificial intelligence is the science and engineering of making intelligent machines, especially intelligent computer programs. Artificial intelligence researchers seek to achieve a scientific understanding of the mechanisms underlying thought and intelligent behaviour and their embodiment in machines. As a result they tend to base their models on cognitive motivation theories. Sloman and Croucher (1981) introduced the need for a "store of 'springs of action' (motives)" as part of a computational architecture of the mind. They presented a broad model of two-tiered control in which motives occupy the top level and provide the drive or urge to produce lower level goals that specify the behaviour of an agent. In a departure from theories of motivation in natural systems, Sloman and Croucher hypothesised that motives would need to incorporate structural descriptions of states to be achieved, preserved or avoided. Many subsequent models of motivation have used motives of this form to trigger goal creation (Aylett et al 2000; Luck and d'Inverno 1998; Norman and Long 1995; Schmill and Cohen 2002).

Norman and Long (1995) proposed a model for **motivated goal creation** that enables agents to create goals both reactively and pro-actively. In their model, the state of the environment and the agent are monitored by a set of motives. Motives are domain specific so, for example, in a warehouse domain motives might include satisfying orders in a timely manner, keeping the warehouse tidy and maintaining security. A change in state may trigger a response from the agent if the strength of the motive exceeds a certain threshold. The strength of a motive is calculated from the state of the domain and the internal state of the agent. This mechanism ensures that the agent will only respond to and reason about changes if they are sufficiently important. A motivational response creates a goal to consider the primary reasons for the trigger of attention. If further activity is required a goal is created that will, if satisfied, cause the mitigation of the stimuli that triggered the motivation. This mitigation takes the form of planned actions in the agent's domain. Because these agents can predict their future beliefs they are able to create proactive goals.

The notion of domain specific goals is a departure from many of the motivation theories from natural systems which seek to find general principles to describe the causation of action. Practically speaking, the implementation effort required to create agents for new domains is high as new motives must be defined for each new domain. In a more domain independent approach, Saunders and Gero created agents motivated by **curiosity** by extending the social force model (Helbing and Molnar 1995) with a fifth force, a desire to move towards potentially interesting physical locations. Various general theories of what is interesting have been developed. For example, Lenat's AM (1976) included 43 heuristics designed to assess what is interesting. Schmidhuber (1997) defined something to be interesting if it is 'similar-yet-different' to previously experienced situations. Saunders and Gero (2002; 2004) drew on these examples to develop computational models of novelty, interest and curiosity. Saunders defined the novelty of environmental stimuli to be inversely proportional to how often stimuli are experienced, how similar stimuli are to each other and how recently stimuli have been experienced. He implemented a novelty detector using the classification error from unsupervised neural networks. Like Schmidhuber, Saunders and Gero believe that novelty is not the only determinant of interest. Interest in a situation is also related to how well an agent can learn the information gained from novel experiences. Consequently, the most interesting experiences are often those that are 'similar-yet-different' to previously encountered experiences. Saunders and Gero in their curious agent, model this phenomonen using the arousal theory of motivation (Berlyne 1960).

While curious agents incorporate a domain independent model of interest, and are able to learn an underlying model of the situations they encounter in their environment, they are unable to learn new behavioural sequences to manipulate their environment. One well known model for learning such behavioural sequences is the **reinforcement learning** algorithm (Sutton and Barto 2000). Reinforcement learning uses rewards to guide agents to learn a function which represents the value of taking a given action in a given state with respect to some task. An agent is connected to its environment via perception and action. On each step of interaction with the environment, the agent receives an input that contains some indication of the current state of the environment. The agent then chooses an action as output. The action changes the state of the environment and the value of this state transition is communicated to the agent through a scalar reinforcement signal. The agent's behaviour should choose actions that tend to increase the long-run sum of values of the reinforcement signal. This behaviour is learnt over time by systematic trial and error. Reinforcement learning parallels operant theory with the additional assumption that reinforcement stimuli are always provided from a source external to the agent.

In an effort to produce agents that are able to bootstrap a broad range of competencies in a wider range of domains, Singh et al (2005) developed a model of **intrinsically motivated reinforcement learning**. In this model, agents are hardwired to identify changes in light and sound intensity as salient (interesting) events. Each first encounter with a salient event initiates the learning of an option and an option-model (Precup et al 1998) for that salient event. An intrinsic reward is generated each time the salient event is encountered that is proportional to the error in the prediction of the salient event according to the learned option-model for that event. When the agent encounters an unpredicted salient event a few times, its updated action-value function drives it to repeatedly attempt to achieve that salient event. The agent acts on the environment according to an e-greedy policy with respect to an action-value function that is learned using a mix of Q-learning and SMDP planning. As the agent moves around the world, all the options and their models initiated so far are simultaneously updated using intra-option learning algorithms. Initially only the primitive actions in the environment are available to the agent. Over time, the agent identifies and learns skills represented by options and option models. These then become available to the agent as action choices. As the agent tries to repeatedly achieve salient events, learning improves both in its policy for doing so and its option-model that predicts the salient event. As the option policy and option model improve, the intrinsic reward diminishes and the agent becomes bored with the associated salient event and moves on.

3.3.3. Motivation Theories Derived from Social Interaction Theories

Social motivation theories have been pursued by researchers in artificial life and artificial intelligence alike and frequently inform the development of multi-agent systems. One such example from the artificial life community experiments with the **evolution of purposeful behaviour** in multi-agent systems. Gusarev et al (2001) experimented with a population of motivated agents capable of reproducing. Their simulation consisted of a population of agents with two basic needs, the need for energy and the need to reproduce. The population evolves in an environment where patches of food grow. Agents can move, eat grass and mate with each other. Mating results in a new agent that inherits the characteristics of its parents according to some simple genetic rules. Their simulation demonstrated that simple hierarchical control systems in which simple reflexes are controlled by motivations, can emerge from evolutionary processes. They showed that this hierarchical system is more effective compared to behavioural control governed by means of simple reflexes only.

In another example from artificial intelligence literature, Martindale (1990) presented an extensive investigation into the role that the search for novelty plays in literature, music, visual arts and architecture. He concluded that the search for novelty exerts a powerful influence on creative activity. Saunders and Gero (2001) produced a computational model of creativity that captures the search for **novelty within a social context**. In his model, agents can communicate particularly interesting artworks to others as well as reward other agents for finding interesting artworks. He shows that both an individual's need for novelty and the collective experience of a group of agents are responsible for creating a consensus as to what is creative.

4. Intelligent Rooms

We take the concept of cognitive agents as a component of elements of a virtual world, extend the agent model to include motivation and curiosity, and move the agent from being associated with an element of a virtual world to being associated with a physical place. This extends recent work in intelligent rooms.

The idea of rooms with embedded computing power has been a subject of research in computer science since at least the late 1960s. Krueger's work on rooms that users can interact with such as VIDEOPLACE (Krueger, 1985) and the work of MIT's Architecture Machine Group on novel user interfaces for rooms using gesture and speech recognition systems such as in "Put-That-There" (Bolt, 1980) laid the groundwork for current research.

More recently, in their papers on Intelligent Environment (IE) design, Brookes et al (1997) and Coen (1998) argued that a key design goal for developing IEs is to enable them to adapt to, and be useful for, everyday activities. They also adopted the stance shared by Ubiquitous Computing researchers that the computing power embedded in IEs ought to be invisible and integrate naturally with everyday activities (Weiser, 1991), in contrast with earlier systems which often required users to be familiar with special interface devices in order to interact with computer-enhanced rooms. Since these papers were published the focus in IEs has been on multi-layer system architectures and sensor hardware appropriate for the IE computational models rather than incorporating creative adaptive capabilities. 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 creatively adapt its behaviour from its patterns of usage and learn for itself how to manipulate its sensors and effectors usefully both of these issues could be addressed.

An agent controlled IE that is a physical space for living or working can bring embedded computational power 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 (2004) 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. The agent that makes the environment intelligent is in the layer of software between low-level hardware management and higher-level application management, able to receive input from the room's sensor systems and to utilize effectors and direct applications.

IEs have several specific design requirements. In addition to adapting to, and being useful for, everyday activities as mentioned above, Brooks et al and Coen have argued that IEs should have a high degree of interactivity and should be able to understand the context in which people are trying to use them and behave appropriately. An IE is essentially, as Kulkarni (2002) 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 Figure 3, attempts to facilitate everyday activities via a system called ReBa, described by Hassens et al (2002) 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 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 are programmed to recognize the actions of the user and take an appropriate action. It does not creatively adapt to new usage patterns. Furthermore, when new sensors are added to the room, the existing rules are modified manually if they are to take advantage of the new sensor information. Our model, by contrast, uses a mechanism driven by a motivation model to creatively adapt its behaviours rather than having the behaviours predefined as part of the agent.

Other researchers have taken approaches to designing environments that are not explicitly agent-based. Both the University of Illinois' Gaia (Roman et al 2002) and Stanford University's Interactive Workspace Project (Hanssens et al 2002) have taken a more OS-based approach, developing Active Spaces and Interactive Workspaces respectively, which focus on the space's role as a platform for running applications and de-emphasizing the role of the room as a pro-active facilitator. The onus for prompting action from the space in these systems is placed upon the user and the applications developer, and although Gaia's context service provides the tools for applications developers to create agent-based facilitating applications, their overall model is not inherently reactive or adaptive. Georgia Tech's Aware Home Research Initiative plans on incorporating an infrastructure for developing context-aware applications (Kidd et al 1999). Our model focuses on the adaptive behaviour of the IE rather than on the hardware/software infrastructure of the sensors and effectors.

Figure 3. MIT's Intelligent Room Project (Hassens et al 2002)

5. An Example of a Design that is Creative: A Curious Place

The requirement for creative adaptability in an IE suggests the need to focus the agent's attention on novel events in the agent's environment in order to develop new behaviours from observations of patterns in these events. From the literature on motivation as a mechanism for focusing attention discussed in the previous section, we present a model driven by intrinsic motivation, effectance motivation, and arousal. From Sing et al's work on a computational model for intrinsic motivation comes the idea of an agent being self-motivated to learn about novel events. Combining these structures with a frequency-based model of curiosity and learning and action components creates a model of a motivated learning agent for an IE, based on a more general model presented by Kasmarik et al (2005), that creatively adapts its own behaviours in response to the recognition of curious events that it senses.

Our curious place model, shown in Figure 4, 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 by the curious place. The agent has sensors to sense the state of the world, effectors to change certain aspects of the state of the world, a memory of world states and events, and a reasoning process that includes motivation, action, and learning.

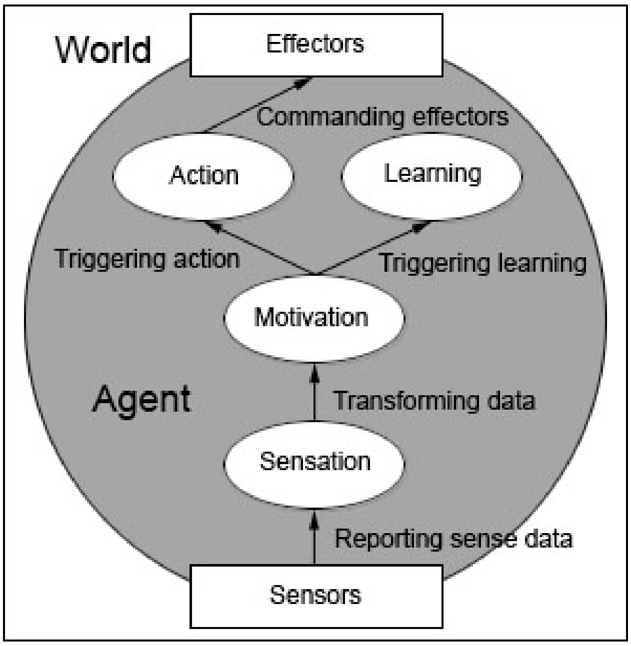

Figure 4. The curious place model.

5.1. THE WORLD STATE

The curious place exists within a specific world available through its sensors. 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 creatively adapting to new behaviour patterns of its users. 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 state of a pressure pad in the floor can only be changed by a person moving on or off the pressure pad, and therefore, cannot be directly changed by the agent. 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 changed by a person or by the agent. A world state at time *t*, is sense data represented in the following form:

```
W(T) ::= ( <senseOnly> | <senseEffect> )
```

And an example of such a state is:

```
(PRESSURE_PAD=ON | LIGHT _INTENSITY=0.5, DESK_LAMP=ON)
```

This distinction is relevant because the agent should only be motivated to learn to repeat events over which it has control. For example, the change in the state of the pressure pads can only be made by humans moving themselves or objects in the room, so an event that includes only a change in the state of the pressure pads cannot be affected by the agent. In contrast, an event that includes a change in the state of the data projectors can be affected by the agent.

5.2. 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 compared with the previous world state *W(t-1)* to extract events. An event is represented as *Delta(t)*, the changes in sensor inputs between *W(t)* and *W(t-1)*. These changes are central to the arousal model of the motivated learning agent, since arousal is a response to a change in stimulus levels. *Delta(t)* takes the same form as *W(t)*, a partitioned tuple, but the values of the tuple represent the changes in sense data values 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_INTENSITY=0.5, DESK_LAMP=ON)
W(1) = (PRESSURE_PAD=OFF | LIGHT_INTENSITY=0.8, DESK_LAMP=ON)
Delta(1) = (PRESSURE_PAD=1 |LIGHT_INTENSITY=0.3, DESK_LAMP=0)
```

5.3. MOTIVATION

The motivation component serves to focus the attention of the agent on novel events and guide it to infer new behaviours from interesting patterns in its history of states. Identifying novel events is governed in this model by a motivational model of curiosity. While a novel event is considered curious to the agent and therefore motivates it to learn new behaviours, the event can't be so novel that there is insufficient data about the world states occurring before the event for the agent to adapt or learn new behaviours. We considered two models that capture this idea of "different, but not too different": the self organising maps used by Saunders and Gero in their curious agent, and the novelty detector based on event frequency clusters used by Kasmarik et al in their motivated agent.

We are currently working with the novelty detector based on event frequency clusters. Events are divided into groups using unsupervised clustering of event frequencies. Each group is defined to be novel or not novel based on their frequencies of occurrence. The novel events are then further clustered into groups of increasing rarity so that the agent can be motivated to learn about more common or 'easier' events that are more likely to have sufficient patterns in the agent's memory.

Clustering is performed by first sorting events in order of ascending frequency where frequency is calculated as the number of times the event has occurred divided by the size of the agent's lifetime. This produces an ordering $(e_1, f_1), (e_2, f_2) \ldots (e_n, f_n)$ with differences $d_1, d_2 \ldots d_n$ where $d_k = f_k - f_{k-1}$. K-means clustering with $k=2$ and initial centroids 0 and d_{max} where $d_{max} = \max_j d_j$ produces two groups g_1 and g_2 with a verage distances to centroids a_1 and a_2. g_i has the minimum average a_i then events can be clustered as follows: Place f_1 in a new cluster. For $f_2, f_3 \ldots f_n$, place f_k in the same cluster as f_{k-1} if $d_k \in g_i$ or in a new cluster otherwise. We say that an event e_i is novel if its frequency f_i falls in the same cluster as f_1.

5.4. LEARNING

The learning component is at the core of the agent's creative adaptability. Because it is inappropriate for an intelligent room to experiment with changes in the state of the room as a reinforcement learning agent might, learning must rely upon drawing inferences from previously experienced world states via data mining techniques without being able to affect 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 delta combined with hashes of recently activated behaviours and `<action>` is a tuple of attribute-value pairs consisting only of sensor data that can be affected. 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. Techniques based on Generalised Sequential Pattern Mining (Agrawal and Srikant 1995; Srikant and Agrawal 1996) or the MINEPI algorithm (Manilla et al 1997) can then be used to find these rules from the memory of world state tuples. The more that a mined rule seems like correct behaviour to the agent in terms of support and confidence, the more strongly the agent feels motivated to follow the rule, mirroring a kind of effectance motivation. This effectance motivation is represented by a priority associated with each mined rule that is a function of the support and confidence assigned to the rule by the data mining algorithm.

5.5. 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 agent's effectors. This is done via a rule engine that supports rule prioritisation. Behavioural rules that generate more effectance motivation are favoured over those that produce lower effectance motivation as reflected in the rule's priority. It then sends the appropriate commands to the agent's effectors to enact the changes in the world dictated by the rule selected.

5.6. MEMORY

The motivation, learning, and action components all require information about earlier states of the world, and the sensation and learning components update that information. This necessitates a memory component for storing deltas and behavioural rules. The various kinds of interactions with memory are outlined in the sections on the specific components and Figure 5 provides a diagrammatic summary.

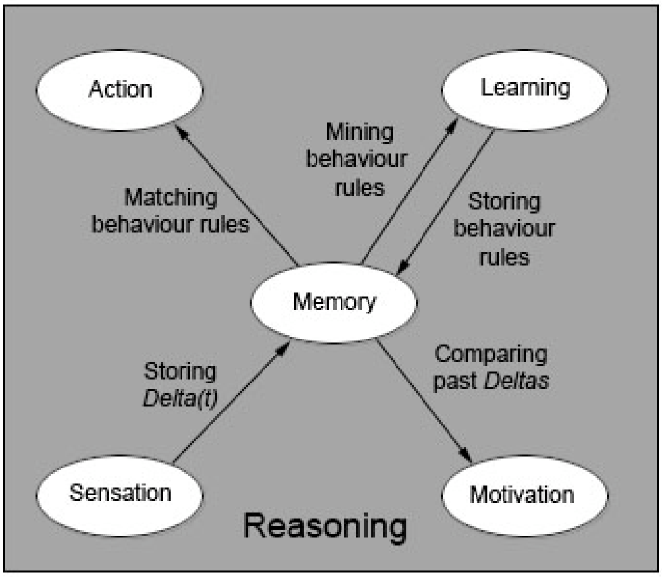

Figure 5. Interactions between the reasoning and memory components of the model.

6. Conclusions

Computational models of creative design are usually considered as tools that assist a human designer while designing. In this paper we reconsider this assumption and look at computational models of creative design as a component of a design product. The situations in which this is relevant are scenarios in which a design product operates in an unpredictable and changing environment, for example, a virtual or physical room that supports a broad range of changing human activities.

We have examined the key characteristics of models of creative design and focussed on a review of motivation theories. Our objective is to establish a theoretical basis for a design product capable of exhibiting creative design behaviour, and specifically a computational model of a curious place. The core of our model is a computational model of motivation that drives an agent for a room to learn new behaviours. The model that we present is currently being implemented and tested in The Sentient, a multipurpose room for seminars, design cognition experiments, and immersive interactive displays.

Our model for a creative design differs from other models of interactive designs by comprising: a model of motivation to trigger learning, a learning component that can use its previous experiences as a basis for developing new behaviours, and an action model that can determine which known behaviour to apply to change the environment in response to a change in the state of the world.

7. References

Agrawal, R, Srikant, R: 1995, Mining Sequential Patterns, *Proc. of the Int'l Conference on Data Engineering (ICDE)*, Taipei, Taiwan, March 1995.

Alderfer, C: 1972, *Existence relatedness and growth*. Free Press, New York.

Atkinson, JW and Feather, NT: 1966, *A theory of achievement motivation*. Wiley, New York.

Atkinson, JW and Raynor, JO: 1974, *Motivation and achievement*. V. H. Winston, Washington D. C.

Avila-Garcia, O and Canamero, L: 2002, Comparison of behaviour selection architectures using viability indicators, *The EPSRC/BBSRC International Workshop on Biologically Inspired Robotics: The Legacy of W. Grey Walter*, HP Labs, Bristol, UK, pp. 86-93.

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

Berlyne, DE: 1970, Novelty, complexity and hedonic value. *Perception and Psychophysics*, **8**: 279-286.

Berlyne, DE: 1971, *Aesthetics and psychobiology*. Prentice-Hall, Englewood Cliffs, NJ.

Boden, M: 1994. What is creativity? In Margaret Boden (ed.) *Dimensions of creativity,* MIT Press.

Bolt, RA: 1981, Put-That-There: Voice and gesture at the graphics interface, in *Proceedings of the 7th Annual Conference on Computer Graphics and Interactive Techniques*, Seattle, US, pp. 262-270.

Brooks, RA: 1997, with contributions from M. Coen, D. Dang, J. DeBonet, J. Kramer, T. Lozano-Perez, J. Mellor, P, Pook, C. Stauffer, L. Stein, M. Torrance and M. Wessler. The Intelligent Room Project. In *Proceedings of the Second International Cognitive Technology Conference (CT'97)*, Aizu, Japan, pp. 271-279.

Canamero, L: 1997, Modelling motivations and emotions as a basis for intelligent behaviour. In: W.L. Johnson (Editor), *The First International Symposium on Autonomous Agents*. ACM Press, New York, NY, pp. 148-155.

Clark, S and Maher, ML: 2005, Learning and Designing in a Virtual Place: Investigating the Role of Place in a Virtual Design Studio, *Proceedings of Ecaade 2005*, Lisbon Portugal.

Coen, MH: 1998, 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 pp. 547 – 554.

Darwin, C: 1859, The origin of species.

Dasgupta, S: 1994, *Creativity in Invention and Design: Computational and Cognitive Explorations of Technological Originality*. Cambridge University Press.

Deci, E and Ryan, R: 1985, *Intrinsic motivation and self-determination in human behaviour*. Plenum Press, New York.

Gero, J and Maher, ML: 1993, Preface in JS Gero and ML Maher (eds) *Modellng Creativity and Knowledge-Based Creative Design*, Lawrence Erlbaum and Assoc. pp 106.

Gershenson, C: 2001, Artificial Societies of Intelligent Agents. BEng Thesis , Fundacion Arturo Rosenblueth, Mexico.

Green, RG, Beatty, WW and Arkin, RM: 1984, *Human motivation: physiological, behavioural and social approaches*. Allyn and Bacon, Inc, Massachussets.

Gu, N and Maher, ML: 2004, A Grammar for the Dynamic Design of Virtual Architecture Using Rational Agents, *International Journal of Architectural Computing*, **4**(1): 489-501.

Gusarev, RV, Burtsev, MS and Red'ko, VG: 2001, Alife model of evolutionary emergence of purposeful adaptive behaviour, *European Conference on Artificial Life*.

Hanssens, N, Kulkarni, A, Tuchinda, R, Horton, T: 2002, Building Agent-Based Intelligent Workspaces. In: *Proceedings of the 3rd International Conference on Internet Computing,*pp. 675-681

Heider, F: 1958, *The psychology of interpersonal relations*. Wiley, New York.

Helbing, D and Molnar, P: 1995, Social force model for pedestrian dynamics. *Physical Review*, **51**: 4282-4286.

Horner, MS: 1974, The measurement and behavioural implications of fear of success in women. *Motivation and Achievement*: 91-117.

Huhns, MN and Stephen, L: 1999, Multiagent Systems and Society of Agents, *in* Weiss, G. (ed.), *Multiagent Systems, A Modern Approach to Distributed Artificial Intelligence*, MIT Press, MA, pp.79-120.

Hull, CL: 1943, *Principles of behaviour*. Appleton-Century-Crofts, New York.

Hull, CL: 1952, *A behaviour system: an introduction to behaviour theory concerning the individual organism*. Yale University Press, New Haven.

Jeng, T: 2004, 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, pp. 579 – 592.

Kandel, ER, Schwarz, JH and Jessell, TM: 1995, *Essentials of neural science and behaviour*. Appleton and Lang, Norwalk.

Kasmarik, K, Uther, W, Maher, ML: 2005, Motivated Agents, *The 19th International Joint Conference on Artificial Intelligence*, Edinburgh, Scotland, pp. 1505-1506

Kidd, C, Orr, R, Abowd, G, Atkeson, C, Essa, I, MacIntyre, B, Mynatt, E, Starner, T, Newstetter, W: 1999, The Aware Home: A Living Laboratory for Ubiquitous Computing Research. In: *Proceedings of the Second International Workshop on Cooperative Buildings*. Pittsburgh, PA

Kreuger, MW: 1985, VIDEOPLACE--An artificial reality, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, San Francisco, US, pp. 35-40s

Kulkarni, A: 2002, 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, pp. 1-5

Lenat, D: 1976, AM: An artificial intelligence approach to discovery in mathematics. PhD Thesis, Stanford University.

Martindale, C: 1990, *The clockwork muse*. Basic Books, New York.

Maher, ML and Gero, JS: 2002, Agent Models of Virtual Worlds, *Proceedings of ACADIA 2002*, Pamona, CA.
Maslow, A: 1954, *Motivation and personality*. Harper, New York.
McFarland, D: 1995, *Animal Behaviour*. Longman, England.
Mannila, H, Toivonen, H, Verkamo, AI: 1997, Discovery of frequent episodes in event sequences. *Data Mining and Knowledge Discovery*, **1**:259-289.
Mook, DG: 1987, *Motivation: the organisation of action*. W. W. Norton and Company, Inc, New York.
Norman, TJ and Long, D: 1995, Goal creation in motivated agents, *Intelligent agents: theories, architectures and languages*. Springer-Verlag, pp. 277-290.
Román, M, Hess, CK, Cerqueira, R, Ranganathan, A, Campbell, RH, Nahrstedt, KN., Gaia A: 2002, Middleware Infrastructure to Enable Active Spaces. In: *IEEE Pervasive Computing* Oct-Dec (2002) pp. 74-83
Rosenman MA, Smith G, Ding L, Marchant D and Maher ML: 2005, Multidisciplinary Design in Virtual Worlds, in *CAADFutures 2005*, Vienna, Austria.
Runco, M and Pritzker, S eds. 1999, *Encyclopedia of Creativity*. San Diege, CA: Academic Press.
Saunders, R and Gero, JS: 2001, The digital clockwork muse: a computational model of aesthetic evolution, *The AISB'01 Symposium on AI and Creativity in Arts and Science*, SSAISB.
Saunders, R and Gero, JS: 2002, Curious agents and situated design evaluations. In: J.S. Gero and F.M.T. Brazier (Editors), *Agents In Design*. Key Centre of Design Computing and Cognition, University of Sydney, pp. 133-149.
Saunders, R and Gero, JS: 2004, Situated design simulations using curious agents. AIEDAM.
Schmidhuber, J: 1997, What's interesting. TR-35-97, Lugano, Switzerland.
Singh, S, Barto, AG, and Chentanex, N: 2005, Intrinsically motivated reinforcement learning, *Advances in Neural Information Processing Systems 17 (NIPS)*, pp. (to appear).
Skinner, BF: 1938, *Behaviour of organisms*. Appleton-Century-Crofts, New York.
Sloman, A and Croucher, M: 1981, Why robots will have emotions, *The 7th International Joint Conference on Artificial Intelligence*, Vancouver, Canada, pp. 197-202.
Smith, G, Maher, ML and Gero, JS: 2004, Towards Designing in Adaptive Worlds, *Computer-Aided Design and Applications* **1**(1-4): 701-708.
Srikant, R, Agrawal, R: 1996, Mining Sequential Patterns: Generalizations and Performance Improvements, to appear in *Proc. of the Fifth Int'l Conference on Extending Database Technulogy (EDBT)*, Avignon, France, March 1996.
Sutton, RS and Barto, AG: 2000, *Reinforcement learning: an introduction*. The MIT Press.
Tolman, EC: 1932, *Purposive behaviour in animals and men*. Century, New York.
Weiser, M: 1991, The computer for the 21st Century, in *Scientific American*, 265(3), pp. 94-104.
White, RW: 1959, Motivation reconsidered: The concept of competence. *Psychological Review*, **66**: 297-333.
Woodworth, RS: 1918, *Dynamic psychology*. Columbia University Press, New York.
Wooldridge, M.: 1999, Intelligent agents, *in* Weiss, G. (ed.), *Multiagent Systems, A Modern Approach to Distributed Artificial Intelligence*, MIT Press, MA, pp.27-78.
Wooldridge, M and Jennings, NR: 1995, Intelligent agents: Theory and practice*, Knowledge Engineering Review* **10**(2): 115–152.
Wundt, W: 1910, *Principles of physiological psychology*. Macmillan, New York.