Achieving Creative Behaviour Using Curious Learning Agents

Mary Lou Maher, Kathryn Merrick, Rob Saunders

Key Centre for Design Computing and Cognition
University of Sydney
[m.maher | kathryn | rob]@arch.usyd.edu.au

Abstract

Creativity is often associated with surprise, novelty, usefulness and value. These characteristics do not assist in the development of a model for intelligent systems to achieve creative behaviour, since they are characteristics that help identify when something or someone has been creative, as post-facto evaluation. Rather, models of creative behaviour for intelligent systems draw on process models such as analogical reasoning and induction, or on principles such as "make the familiar strange" or "make the strange familiar". This paper describes how a computational model of curiosity, based on cognitive models of novelty and interest, can be used to focus attention in learning agents. We show how this combination of curiosity and learning can be the core reasoning process in agent-based systems that achieve creative behaviour.

Creative Behaviour

Psychologists and biologists have proposed a number of definitions for creative behaviour in natural systems. In humans, creative behaviour may be defined as behaviour that results in a product that is unique or valuable to either an individual or a society. Alternatively, from a behaviouristic viewpoint, creative behaviour may be defined as a unique response or pattern of responses to an internal or external discriminative stimulus (Razik, 1976). In animals, the definition of creative behaviour is somewhat weaker, with creative behaviour considered as behaviour that does not occur in the normal activity of a given species (Pryor et al., 1969).

The design of artificial systems that exhibit creative behaviour has been pursued with the goals of achieving systems that are autonomous, pro-active and able to find novel solutions to complex problems. Creative behaviour that results in creative products, for example, is of particular interest to designers for the development of powerful design tools (Gero, 1992; Saunders, 2001).

This paper considers creativity from a behaviouristic viewpoint. That is, we consider the design of artificial systems that can generate unique responses or patterns of responses to internal and external stimuli. Artificial systems capable of this kind of creative behaviour are particularly relevant to applications that require creative expression through action rather than applications that aid in the design of creative products. Developing creative behaviour in artificial systems focuses on the automatic generation of sequences of actions that are novel and useful, in contrast to artificial systems that model the automatic generation of creative solutions based on cognitive models of creativity such as analogy (Goel, 1997; Maher and Balachandran, 1994), evolutionary systems (Maher, 1994; Alem and Maher, 1994), situated reasoning (Gero, 2006) or flexible ontology (Gero and Kannengiesser, 2007). Example applications for artificial systems that exhibit creative behaviours include non-player characters in computer games, virtual actors in animated films, intelligent environments that adapt to new patterns of usage, or virtual assistants that can proactively identify and act to solve tasks.

To achieve creative behaviour in artificial systems, we draw on our understanding of creative behaviour in natural systems for inspiration. We begin in the next section with a review of literature describing processes that are thought to contribute to creative behaviour in natural systems. Following this, we describe two approaches to artificial systems that achieve creative behaviour by generating behavioural diversity through attention focus on novel or interesting events. The first approach generates emergent creative behaviour with an intrinsic reward function for reinforcement learning based on models of novelty and interest. We demonstrate this model for the control of non-player characters in games. The second approach generates emergent creative behaviour using a similar model of novelty and interest to focus on specific external stimuli in supervised learning. This model is demonstrated as a virtual assistant that controls devices in a virtual meeting room.

Generating Creative Behaviour

As with the definition of creative behaviour, there exists a range of theories regarding the processes that generate creative behaviour (Gorney, 2007). In the early 1900s, Freud (1900) explained creativity as a process of reducing the tension between fundamental biological drives, social norms and restrictions. In contrast, Maslow (1968) later believed that creativity was motivated by a cognitive need for self-actualisation. He described creative behaviour as a process of spontaneous expression by a person whose more basic biological needs have been satisfied.

Kirton's (1994) work draws a number of alternative explanations for creativity together in the Kirton Adaptive-Innovative Inventory. Kirton (1994) categorises creative behaviour under two headings: adaptive creative behaviour and innovative creative behaviour. Adaptive creative behaviour describes the dynamic interplay between a person and their environment that is necessary for their survival. Innovative creative behaviour can be thought of as satisfying the cognitive search for individuality,

meaning (Storr, 1989) or self-actualisation (Maslow, 1968).

The idea that there may be different types of creativity is also captured through the distinction between extrinsically and intrinsically motivated creativity (Amabile, 1983; Amabile and Collins, 1999). Extrinsic motivation is characterised by a focus on external reward, recognition or direction of an individual's behaviour. In contrast, intrinsic motivation prompts an individual to engage in activity primarily for its own sake, because the individual perceives the activity as interesting, involving, satisfying or personally challenging. Amabile's (1996) intrinsic motivation principle states that intrinsic motivation is conducive to creativity while controlling extrinsic motivation can be detrimental. However, informational or enabling extrinsic motivation can be conducive to creative behaviour, particularly if there is also a high level of intrinsic motivation.

Boden's (1994) work also provides insight into intrinsic and extrinsic creativity. Boden (1994) proposes two types of creativity: psychological creativity within an individual and historical creativity within a society. Csikszentmihalyi (1996) goes further to propose specific types of intrinsic motivation that are conducive to creative behaviour. He believes that typical motivation for creativity is a combination of personal interest and a sense of discordance in the environment. The creative process is thus a search for interest by changing the environment to reduce the discordance. Similar views are held by Martindale (1990), who proposes that the search for novelty is a key motivation for creative individuals. Creative individuals progressively change the way they interact with their environment to attain novelty and avoid replication. Behavioural diversity is thus a key factor in attaining creative behaviour, as is the ability to introduce new variables and patterns progressively into the creative process (Gero, 1992).

One way in which new patterns and variables are introduced into a structure or process is through the phenomenon of emergence (Gero, 1992). Emergence can occur in structure, function or behaviour. The idea of emergent behaviour in an artificial system describes patterns of action that cannot be directly traced back to the system's components, but rather emerge as a result of the way those components interact (Steels, 1990). When emergent behaviour is novel, unique or interesting is added to an artificial agent's repertoire of behaviours, for example through learning, the agent can be viewed as displaying creative behaviour.

In summary, creative behaviour can be seen as a product of motivation – both extrinsic and intrinsic – adaptation, novelty, behavioural diversity, emergence and learning. There are numerous other characteristics that are also associated with creativity, such as unexpectedness, uncommonness, peer recognition, influence, intelligence and popularity (Runco and Pritzker, 1999). However, we draw on the former set of concepts as a starting point for developing and evaluating models of creative behaviour in the next sections.

Towards Creative Behaviour using Curious Learning Agents

Two important, positively recurrent traits of creative people are autonomy (Feist, 1999) and curiosity (Davis, 1999). Saunders (2001) focused on the role of curiosity in creativity to develop computational models of creativity as a search for novelty and interest. In Saunder's (2001) model, novelty $N_{(t)}$ is computed using a real-time novelty detector (Marsland et al., 2000). This model combines a self-organising map (SOM) with a layer of habituating neurons, which respond to activity in the SOM using Stanley's (1976) model of habituation. According to Stanley's (1976) model, the novelty of any given stimulus fluctuates over time, though cycles of habituation and recovery. This is shown in Figure 1.

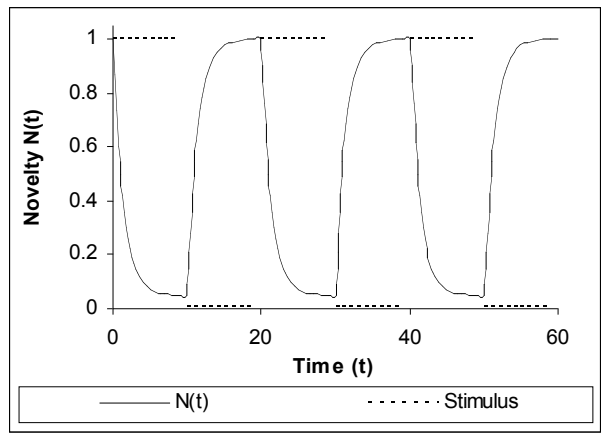

Figure 1. The novelty of a stimulus fluctuates over time. From (Merrick, 2007b).

Saunders modifies novelty using the Wundt (1910) curve to compute interest. He uses two sigmoid functions to provide positive feedback as novelty increases at low levels, and negative feedback for very high novelty. Curiosity (or interest) is highest for moderate novelty values as shown in Figure 2. In this paper we use the Merrick and Maher (2006) approximation of the Wundt curve to model curiosity:

$$C = \frac{1}{1 + e^{-10^{+}(2N_{(t)} - 0.5)}} - \frac{1}{1 + e^{-10^{-}(2N_{(t)} - 1.5)}} \qquad (1)$$

Saunders and Gero (2001; 2002) developed models of curiosity in the context of curious flocking agents and in a social context based on Martindale's (1990) thought experiments. Later work has combined models of motivation, including curiosity, with learning to achieve autonomous mental development in artificial agents (Merrick and Maher, 2008; Oudeyer et al., 2007). This

combination of curiosity and autonomous learning suggests a starting point for modelling creative behaviour.

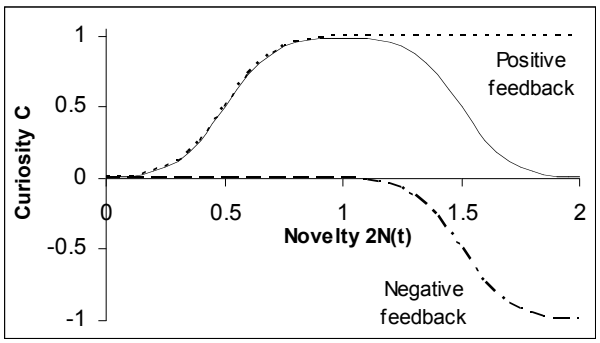

Figure 2. Interest is highest for moderate novelty values. From (Merrick and Maher, 2006).

Merrick and Maher (2008) presented two general models of motivated learning for the design of adaptive reinforcement learning (Sutton and Barto, 2000) or supervised learning (Nilsson, 1996) agents. Merrick and Maher (2008) focus on achieving adaptive behaviour in a reinforcement learning setting. In this section we describe two specific models for curious learning agents using motivated reinforcement learning and motivated supervised learning, which focus, instead, on achieving different types of creative behaviour for different applications.

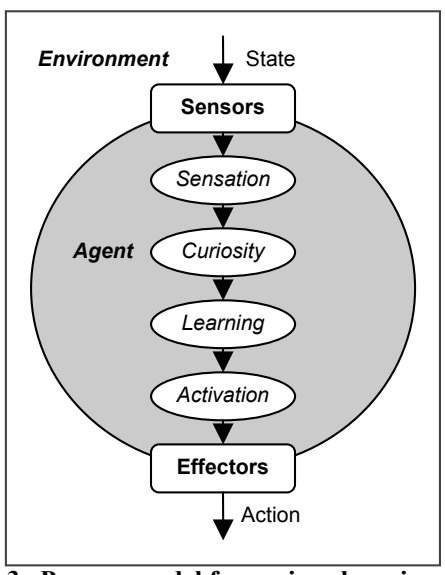

Figure 3. Process model for curious learning agents.

A process model for our curious learning agent is shown in Figure 3. The sensation process transforms raw data about the environment into structures to facilitate further reasoning. Structures include observations of the environment, the change or 'event' between the current and previous observed states and example actions performed by humans or other agents. Observations or events (Merrick and Maher, 2008) become stimuli for the curiosity process. The curiosity process reasons about current stimuli by computing novelty and interest. The learning process performs a learning update to incorporate observations, examples and actions into a policy defining how the agent should act. Finally, the activation process selects an action to perform from the learned policy.

The following sections trace two paths through this model for the purpose of designing creative agents as non-player characters in games and as a virtual assistant to control devices in a meeting room.

Creative Behaviour by Experimenting

This model describes a curious learning agent that is intrinsically motivated to explore and experiment in its environment. It adapts its behaviour using reinforcement learning to learn to repeat interesting events, using the algorithm summarised in Figure 4. The curious reinforcement learning algorithm uses an incremental model that cycles through the sensation, curiosity, learning and activation processes. This structure is implemented as a main loop in Line 1. In Line 2, the agent senses the state $S_{(t)}$ of the environment. We assume an attribute-based representation of the environment, in which attributes s_L take numeric values:

$$S_{(t)} = (s_{1(t)}, s_{2(t)}, s_{3(t)}, \ldots s_{L(t)} \ldots) \qquad (2)$$

We use the variable length state space model described by Merrick and Maher (2007a). Each attribute is identified by a label L that can be used to compare the value of the attribute at different moments in time. This model supports creative behaviour by permitting the introduction of new variables into the state space while the agent is learning.

An event $E_{(t)}$ is computed in Line 3, representing the change between the current state and the previous state:

$$E_{(t)} = (s_{1(t)} - s_{1(t-1)}, s_{2(t)} - s_{2(t-1)}, s_{3(t)} - s_{3(t-1)} \ldots)$$

Like the sensed states from which they are computed, events may also vary in length. For the purpose of computing events, attributes that are not present at either time t or time t-1 are assigned a value zero at that time. The resulting event vector is normalised for use to compute curiosity in Line 4, according to Equation 1.

The curiosity value is used to update a policy π mapping states to actions and curiosity values in Line 5. The agent acts in Line 7 by selecting an action from π with the highest curiosity value for the current state most of the time and a random action some small percentage of the time. In the experiments in this paper, we used table-based Q-learning to update the policy π (Watkins and Dayan, 1992). The curiosity value is effectively the reward signal directing learning.

```
1.  Repeat (forever):
2.     Sense S(t)
3.     Compute E(t)
4.     Compute C(E(t))
5.     Update π with S(t-1), A(t-1), C(E(t))
6.     S(t-1) ← S(t); A(t-1) ← A(t)
7.     Execute A(t) from π
```

Figure 4. Curious reinforcement learning.

Creative Behaviour by Observing

This model describes a curious learning agent that is motivated by both internal and external observations to act in its environment. It adapts its behaviour using supervised learning to learn to repeat interesting examples.

This algorithm is summarised in Merrick and Maher (2008). Once again, the algorithm takes an incremental approach to learning, controlled by the main loop in Line 1. In Line 2, the agent senses either the state $S_{(t)}$ of the environment or an example $X_{(t)} = (S_{(t)}, A_{(t-1)})$ of human behaviour. States again use the variable length, attribute-based representation in Equation 2, while actions are enumerated. When a state is sensed, an observation of that state is computed in Line 3 as:

$$O_{(t)} = (s_{1(t)}, s_{2(t)}, s_{3(t)} \ldots)$$

When an example is sensed, an observation is computed combining the sensed state and action:

$$O_{(t)} = (s_{1(t)}, s_{2(t)}, s_{3(t)} \ldots A_{(t)})$$

The curiosity of the observation is then computed in Line 4. First a numeric value $\mathcal{C}(O_{(t)})$ is computed using Equation 1. Then a step function is then used to determine whether the curiosity is high enough to motivate learning and/or action as shown in Figure 5.

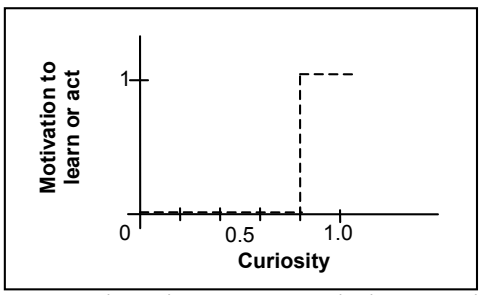

Figure 5. Relationship between curiosity, learning and activation in curious supervised learning agents.

If the agent senses an example and curiosity is greater than a threshold C, then a policy π mapping states to actions is updated by strengthening the weight between the state and action in the current example as in Line 6. The weight between the state and any other action is weakened. If the agent sensed a state and curiosity is above the threshold, the agent acts by selecting the state with the highest weight from the policy π in Line 8. In the experiments in this paper, we used association-based supervised learning to update and select from the policy π (Steels, 1996).

```
1.  Repeat (forever):
2.    Sense S(t) or X(t)
3.    Compute O(t)
4.    Compute step(C(O(t)))
5.    if (sensed X(t) & step(C(O(t)))=1):
6.       Update π with X(t)
7.    if (sensed S(t) & step(C(O(t)))=1):
8.       Execute A(t) from π
9.    S(t-1) ← S(t); A(t-1) ← A(t)
```

Figure 6. Curious supervised learning.

Unlike standard supervised learning, curious supervised learning focuses on a subset of examples that changes progressively over time. Likewise, action occurs in a changing subset of states, depending on the agent's current focus of interest. This allows the agent to selectively focus on or ignore certain stimuli in its search for creative behaviour.

In curious supervised learning, unlike curious reinforcement learning, learning and action are independent so the agent may learn based on examples concerned with one set of states, but act in a different set of states. This means that there is potential for different models of curiosity or other motivating forces to be used to direct learning and action. In this paper we use separate instances of the same curiosity model to direct both learning and activation.

Evaluating Creative Behaviour

Evaluating whether the behaviour of an artificial system is creative is a difficult task. The behaviour of the system cannot be validated using the principles that underlie the approach, yet these principles are important indicators of creative behaviour.

In this paper, we begin by evaluating a weak definition of creativity using principles of creativity in animals (Pryor et al., 1969). In particular, we focus on evaluating the emergence of new behaviours or patterns of behaviour that have not been seen previously. We use a number of quantitative metrics to characterise this behaviour as well as studying the behaviour of curious learning agents qualitatively in different domains.

From a quantitative view, we study the emergence of new patterns of behaviour. The emergence of patterns implies development of behavioural diversity that avoids replication of action. Martindale considered this an important aspect of creativity in natural systems. In this paper, we count emergent, learned behaviour that are repeated at least five times in response to identical stimuli.

From a qualitative perspective, we discuss the emergent behaviour in terms of its focus, novelty and unexpectedness.

Curious Learning Agents as Non-Player Characters

Artificial agents that exhibit creative behaviour are of particular relevance to the computer games industry. The design of non-player characters that display creativity in their behaviour opens the way for new types of games in which characters adapt in surprising and interesting ways to the actions of player characters. This can extend the lifetime of the game by offering players a unique experience each time they play (Merrick, 2007b).

In this section, we analyse the behaviour of sheep in a simulation game. In the game, players are asked to build and script objects to attract the attention of six sheep. The

sheep respond using curiosity and learning. One object built by a player, Sahi Kipling, is the food machine shown in Figure 7. The food machine has a button that, when pressed by a sheep, releases food down the shoot and into the trough. The food is edible and disappears when eaten. Examples of sensed states in this environment, using the labelled, attribute-based format are:

$$S_{(t)} = (1_{foodMachine}, 2_{food}, 1_{sahiiKipling})$$
$$S_{(t+1)} = (1_{foodMachine}, 3_{food}, 1_{sahiKipling})$$

Given this sequence of states, the event that occurs at time t+1 is:

$$E_{(t+1)} = (+1_{food})$$

Over the course of two hours, the six sheep engaged in five different behaviours as follows:

- Press the button, move to the food, eat the food, move back to the button…
- Press the button twice, move to the food, eat both the food balls, move back to the button…
- Press the button three times, move to the food, eat the three food balls, move back to the button…
- Move to the player, move to the food machine, move back to the player…
- Move to the food shoot, press the button, eat the food, press the button, eat the food…

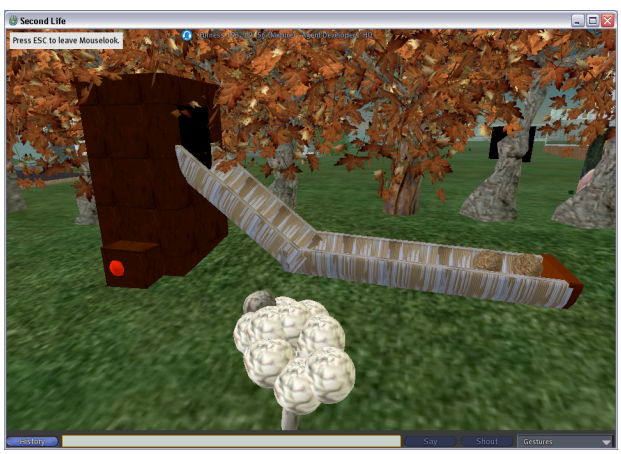

Figure 7. A non-player character (sheep) controlled by a curious reinforcement learning agent interacts with a food machine.

Anecdotally, the diversity of behaviour was surprising because the food machine was designed with only the first behaviour in mind. The fifth behaviour was particularly creative, with the sheep wedging itself on the food shoot so that it could reach the button and allow the food to roll into its mouth without having to move. The fourth behaviour was performed by a sheep that due to its experiences with other player built objects, did not find the food machine interesting. The fact that none of the sheep developed behaviours that involved pressing the button more than three times suggests a limitation of our current approach, in that it does not appear to support the emergence of behavioural patterns of longer sequences of actions. The design of models of curiosity that will support more complex behaviour is an open research problem.

The combination of curiosity and reinforcement learning in this example produces a kind of creativity that is intrinsically motivated and may be either psychologically creative or historically creative within the flock (society) of sheep.

The sheep initially act randomly to explore their environment. During this exploration they identify interesting events and generate internal reward for the learning process. Each new interesting event encountered will thus result in a perturbation of existing behavioural patterns as the learning process adapts to the changing reward signal. Agents are guided towards creative behaviour by the identification of interesting events, but the behaviour is learned after the interesting event is recognised. This is in contrast to existing 'generate-and-test' approaches, such as genetic algorithms (Goldberg, 1989), where behaviour is generated first then evaluated for creativity. The advantage of guiding the system ante-facto, by biasing exploration towards interesting events, lies in the fact that large numbers of behaviours do not have to be generated for evaluation.

The next section considers agents that model a different kind of creative behaviour based on observation of external stimuli.

Curious Learning Agents as Virtual Room Assistants

Agents with creative behaviour also offer an alternative solution to the design of intelligent environments. Previously, the focus of intelligent environment research has been on the development of adaptive middleware architectures for management of, and communication between, resources. The design of behavioural architectures that can explore the potential of new devices and develop behaviours to support human activities has so far received less focus.

In this section, we analyse the behaviour of a curious supervised learning agent controlling devices in the virtual meeting room shown in Figure 8. The agent can monitor the presence of avatars in the centre of the room (away from the doors), lighting, whether avatars sitting in chairs, smart board status, projector screen and teleconferencing software. It can control the lights, smart board, the slide show on the projector and teleconferencing software on the smart board. Users of the virtual meeting room can control the same set of devices. We simulated two users who use the room for different activities: Adel Andrews makes slide show presentations and Sahi Kipling does teleconferencing. Two possible sensed states in this environment are:

$$S_{(t)} = (7_{onLight}, 1_{occupiedChair}, 2_{emptyChair}, 1_{SahiKipling}, 1_{offSmartBoard}, 1_{offPojectorr})$$

$$S_{(t+1)} = (7_{onLight}, 1_{occupiedChair}, 2_{emptyChair}, 1_{SahiKipling}, 1_{onSmartBoard}, 1_{offPojectorr})$$

This sequence of states may have been caused by the following example in which Sahi Kipling executes Action 3 to turn on the smart board:

$$O_{(t)} = (7_{onLight}, 1_{occupiedChair}, 2_{emptyChair}, 1_{SahiKipling}, 1_{offSmartBoard}, 1_{offPojectorr}, 3_{turnOnSmartBoard})$$

A curious supervised learning agent observing activities in the room learned the following five behaviours:

- When Adel is the only avatar sensed, turn on the lights
- When Adel and Sahi are sensed, turn off the lights and turn on the presentation
- When no avatars are sensed, turn off the lights
- When Sahi is the only avatar sensed, turn on the lights
- When Sahi is sitting, launch teleconferencing on the smart board and turn off the lights

Figure 8. Devices in a virtual meeting room controlled by a curious supervised learning agent.

The combination of curiosity and supervised learning in this example produces a kind of creativity that is intrinsically motivated but responds to observations of external activities. The resulting behaviour of the curious supervised learning agent is psychologically creative to the agent, but not necessarily historically creative within the larger society of avatars using the virtual room. This model provides a way for an artificial agent to identify creative behaviours performed by others and learn to support those behaviours.

In this environment, the curious supervised learning agent focused on and learned to mimic only the most frequent human behaviours. Actions performed by avatars less often, although still predictably, tended to be ignored by the curious agent. One possible interpretation of this is that, by focusing on frequently occurring stimuli, curious agents are focusing on important and potentially useful behaviours. In contrast, standard supervised learning will eventually learn any predictable, repeated behaviour, regardless of how irregularly examples are presented. The idea of curious supervised learning is to provide a filter for states and examples to focus learning on important or useful behaviours in environments where it is infeasible for the agent to learn everything.

The tendency of curious agents to focus learning on the most frequently observed events or examples is brought about by a property of the novelty function. The change in novelty (and thus interest) is a function, not only of the number of times a stimulus is encountered, but of the number of significantly different stimuli that occur between encounters. This is true in both the supervised learning and reinforcement learning variants of curious learning agents. This suggests that issue will arise regarding the selection of parameter values of the novelty and interest functions in any given environment to ensure that structured behaviour emerges. For example in an environment with many possible actions, event or examples might occur predictably, but with less frequency than in the experiments in this paper. The parameters of the interest and novelty functions would have to be modified to lengthen the novelty-habituation cycle.

Discussion and Conclusion

The results in this paper show the capacity of curious learning agents to learn structured patterns of behaviour by using curiosity to focus learning on different environmental stimuli. However, the experiments in this paper are conducted in very simple microworlds. Results suggest that the ability of these agents to develop creative behaviour may be limited to relatively short sequences of actions. Moving beyond microworlds to generate behaviour that adheres to a stronger definition of creativity is likely to require different models of curiosity.

In this paper we evaluate creative behaviour in terms of the emergence of behavioural diversity, novelty and unexpectedness. We briefly consider the usefulness of learned behaviour, but in practice, this is particularly difficult to calculate in a general sense, although it may be calculated from surveys of human responses in domain specific cases.

The work in this paper models a phenomenological response to stimuli, in contrast to existing work that generates creative solutions based on existing knowledge. The significance of this work lies in the development of agent models that seek to achieve creative behaviour, as opposed to creative products, and that do this by a process of curiosity-biased exploration and learning.

Acknowledgements

The work was supported by a grant from the Australian Research Council for the project: “Curious Places: Agent Mediated Self-Aware Worlds”. This paper was written while Mary Lou Maher was working for the National

Science Foundation in the USA. Any opinions, findings recommendations or conclusions expressed in this paper are those of the authors and do not necessarily represent the views of the National Science Foundation.

References

Alem, L. and Maher, M.L. (1994). A model of creative design using a genetic metaphor, in T. Dartnall (ed.) *Artificial Intelligence and Creativity: An Interdisciplinary Approach*, Kluwer, pp 281-291.

Amabile, T.M., 1983. *The social psychology of creativity*. Springer-Verlag, New York.

Amabile, T.M., 1996. *Creativity in context: update to the social psychology of creativity*. Oxford: Westview Press.

Amabile, T.M. and Collins, M.A., 1999. Motivation and creativity. In: R.J. Sternberg (Editor), *Handbook of creativity*. Cambridge University Press, Cambridge, pp. 197-312.

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

Csikszentmihalyi, M., 1996. *Creativity: Flow and the Psychology of Discovery and Invention*. HarperCollins Publisher, New York, NY.

Davis, G., 1999. Barriers to creativity and creative attitudes. In: M. Runco and S. Pritzker (Editors), *Encyclopaedia of Creativity*. London: Academic Press, San Diego, CA, pp. 165-74.

Feist, G., 1999. The influence of personality on artistic and scientific creativity. In: R.J. Sternberg (Editor), *Handbook of creativity*. Cambridge University Press, Cambridge, pp. 273-296.

Freud, S., 1900. *The interpretation of dreams*.

Gero, J.S., 1992. Creativity, emergence and evolution in design, *Second International Roundtable Conference on Computational Models of Creative Design*, Sydney, Australia, pp. 1-28.

Gero, JS (2006) Understanding situated design computing: Newton, Mach, Einstein and quantum mechanics, in IFC Smith (ed), *Intelligent Computing in Engineering and Architecture* , Springer, Berlin, pp. 285-297.

Gero, JS and Kannengiesser, U (2007) An ontological model of emergent design in software engineering, *ICED07*, Ecole Centrale de Paris, pp. 70:1-12.

Goel, A., 1997. Design, Analogy and Creativity. IEEE *Intelligent Systems*, 12(3): 62-70.

Goldberg, D., 1989. *Genetic algorithms*. Addison-Wesley.

Gorney, E., 2007. *Dictionary of creativity: terms, concepts, theories and findings in creativity research*. http://creativity.netslova.ru/ (Accessed December, 2007).

Kirton, M.J., 1994. *Adaptors and innovators: styles of creativity and problem solving*. Routledge, London.

Maher, M.L. (1994). Creative design using a genetic algorithm, *Computing in Civil Engineering*, ASCE, pp 2014-2021.

Maher, M.L. and Balachandran, B. (1994). Flexible retrieval strategies for case-based design, Artificial Intelligence in Design'94, J. Gero (ed.), Kluwer Academic Press, pp 163-180.

Marsland, S., Nehmzow, U. and Shapiro, J., 2000. A real-time novelty detector for a mobile robot, EUREL *European Advanced Robotics Systems Masterclass and Conference.*

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

Maslow, A., 1968. *Toward a psychology of being*. Van Nostrand, Reinhold, New York.

Merrick, K., 2007a. Modelling Motivation for Adaptive Non-Player Characters in Dynamic Computer Game Worlds. *ACM Computers in Entertainment*: (to appear).

Merrick, K., 2007b. *Modelling motivation for experience based attention focus in reinforcement learning*. PhD Thesis, School of Information Technologies, University of Sydney.

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

Merrick, K. and Maher, M.-L., 2008. Motivated learning from interesting events: adaptive, multitask learning agents for complex environments. *Adaptive Behaviour*: (to appear).

Nilsson, N.J., 1996. *Introduction to machine learning*. http://ai.stanford.edu/people/nilsson/mlbook.html (accessed January, 2006).

Oudeyer, P.-Y., Kaplan, F. and Hafner, V., 2007. Intrinsic Motivation Systems for Autonomous Mental Development. *IEEE Transactions on Evolutionary Computation*, 11(2): 265-286.

Pryor, K., Haag, R. and O'Reilly, J., 1969. The creative porpoise: training for novel behaviour. *Journal of the Experimental Analysis of Behaviour*, 12(4): 653-661.

Razik, T., 1976. Programming creative behaviour. *British Journal of Education Technology*, 3(7): 5-21.

Runco, M. and Pritzker, S., 1999. *Encyclopaedia of creativity*. Academic Press, San Diego, CA.

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

Saunders, R. and Gero, J.S., 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, J.S., 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.

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

Steels, L., 1990. Towards a theory of emergent functionality, *From Animals to Animats: First International Conference on Simulation of Adaptive Behaviour*. Bradford Books (MIT Press), pp. 451-461.

Steels, L., 1996. Self-organising vocabularies, Technical report, Artificial Intelligence Laboratory Vrije Universiteit Brussels.

Storr, A., 1989. *Solitude*. Free Press, New York.

Sutton, R.S. and Barto, A.G., 2000. *Reinforcement learning: an introduction*. The MIT Press.

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

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