ARTIFICIAL CREATIVITY: A SYNTHETIC APPROACH TO THE STUDY OF CREATIVE BEHAVIOUR

ROB SAUNDERS AND JOHN S. GERO
The Key Centre of Design Computing and Cognition
Department of Architectural and Design Science
University of Sydney, NSW, 2006, Australia
e-mail: {rob, john}@arch.usyd.edu.au

Abstract

We present a novel approach to the computational study of creativity, called *artificial creativity*. Artificial creativity promotes the study of the creative behaviour of individuals and societies in artificial societies of agents. It is similar to the approach to that taken by Artificial Life researchers involved in developing computational models. We present a framework for developing artificial creativity systems as an adaptation of Liu's dual generate-and-test model of creativity. An example implementation of an artificial creativity system is presented to illustrate the potential benefits of our new approach as a way of investigating the emergent nature of creativity in societies of communicating agents. Finally, we discuss some future research directions that are possible by extending the abilities of individuals and studying the emergent behaviour of societies.

1. INTRODUCTION

The aim of artificial creativity is to gain a better understanding of *creativity-as-it-is* in the context of *creativity-as-it-could-be*. In other words, it is the comparative study of creativity as it is found in human societies against creativity as it may be found in artificial societies of agents that may follow quite different social conventions. In this way, the study of artificial creativity is similar to the study Artificial Life; both are synthetic approaches to understanding a complex, and ill-defined behavioural phenomenon, i.e. creativity and life respectively.

The artificial creativity approach provides an opportunity for researchers to study the emergence of creative behaviour in controllable environments, affording a number of possible studies not possible in the real world. The parameters that control the behaviour of individuals can be experimented with to study the affect that they have on the emergence of social structures. In addition, the environment that the society of agents is situated in, e.g. economic conventions, can be adjusted to study the affects that external factors have on the creativity of individuals and societies.

As with Artificial Life, one of the most interesting possibilities of artificial creativity is to be able to re-run history with different starting conditions to find out how products of creative individuals and the structures of creative societies might have differed. For example, by re-running an artificial creativity simulation with different communication policies we can simulate the affect that different communication technologies might have on the development and dissemination of creative ideas.

Artificial creativity is compatible with previous approaches to studying creativity that have developed computational models of creative thinking by allowing them to be deployed within the context of artificial societies as long as they can be embedded within agents that conform to the requirements of artificial creativity. The study of the behaviour of creative thinking within artificial societies provides the opportunity to develop a better understanding of the situatedness of creative processes within socio-cultural situations. As Simon (1981) notes, much of the complexity of human behaviour may come from the complex nature of the environment that they interact with.

The following section provides some background regarding the study of creativity that has lead to the development of the artificial creativity approach described in Section 3. In Section 4, a framework for artificial creativity systems is developed by adapting Liu's dual generate-and-test model. Section 5 presents an example implementation of an artificial creativity system that has been used to computationally investigate some of the predictions made by Martindale (1990) regarding the nature of creativity in societies that value novelty. The results of these experiments are given in Section 6. Finally, Section 7 discusses some future directions for research using artificial creativity systems.

2. CREATIVITY

The need to define the nature of creativity has haunted attempts to develop theories of the processes involved in creative thinking. The difficulty of this task is apparent from the number of definitions that can be found in the literature: Taylor (1988), for example, gives some 50 definitions. Expressed in the definitions of creativity are some widely different opinions about what it means for a person to be creative. From reading the literature, it seems that no agreement may be reached on details of the creative process; however, the definitions provided can be divided into two broad categories.

Firstly, there are definitions of creativity that emphasise creative thinking and promote the view that creativity can be studied solely as a mental phenomenon. These definitions have been a popular in various approaches to studying creativity that deal with individuals, for example, in psychology, cognitive science and artificial intelligence. The models of creativity proposed by Koestler (1964), Newell et al. (1962), and Hofstadter (1979) go into great detail about the cognitive processes involved in creative thinking, particularly the processes involved in the generation of potentially creative ideas. Many of the computational models of creativity are either based directly on these models (e.g. Langley et al., 1987; Hofstadter et al., 1995) or are based on similar models of creative thinking from psychology (e.g. Partridge and Rowe, 1994).

Definitions of creativity in the second category recognise that creativity goes beyond the generation of novel ideas and that society, as the audience of the creative individual, plays an important role in defining what is creative. Creativity is therefore defined with a strong honorific sense that is as much the result of an audience's appreciation of a work as it is the creator's production. Proponents of these definitions contend that creativity cannot occur in a vacuum and must be studied in the context of the socio-cultural environment of the creator (Csikszentmihalyi, 1988; 1999). This definition has been popular in fields that consider the creativity of multiple individuals over extended periods of time, for example, in history, sociology and anthropology (e.g. Martindale, 1990).

Some researchers have attempted to combine these two views of creativity into unified theoretical frameworks. However, the resulting frameworks often maintain the distinction between personal and socio-cultural notions of creativity, as in Boden's P-creativity and H-creativity (Boden, 1990) and Gardner's small-c and big-c creativity (Gardner, 1993).

2.1.1.A Systems View of Creativity

When Csikszentmihalyi developed his systems view of creativity, he turned his attention away from the question "What is creativity?" and focussed upon the issues surrounding the question "Where is creativity?" Importantly, Csikszentmihalyi questioned the mentalistic assumption that creative processes are only to be found in the mind of the creative individual. Instead he proposed that processes essential to creativity, whether personal or socio-culturally defined, are to be found in the interactions between individuals and the society that they are situated within.

The systems view of creativity was developed by Csikszentmihalyi as a model of the dynamic behaviour of creative systems that include interactions between the major components of a creative society (Csikszentmihalyi, 1988). Csikszentmihalyi identified three important components of a creative system; firstly there is the *individual*, secondly there is a cultural, or symbolic, component called the *domain*, and thirdly there is a social, or interactive, component called the *field*. A map of the systems view of creativity is presented in Figure 1.

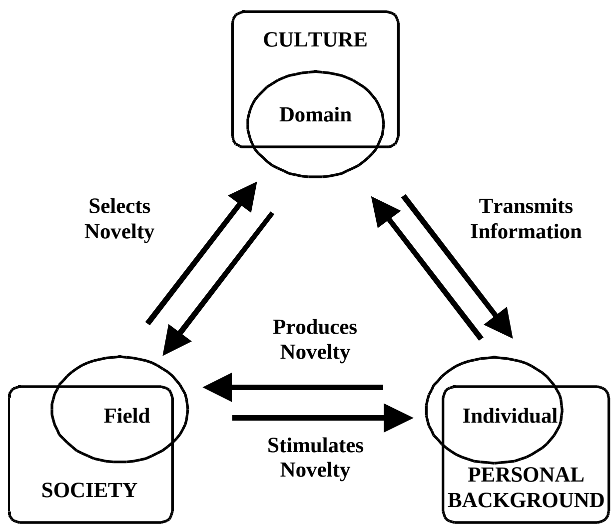

Figure 1: Csikszentmihalyi's systems view of creativity (after Csikszentmihalyi, 1999).

An individual's role in the systems view is to bring about some transformation of the knowledge held in the domain. The field is a set of social institutions that selects from the variations produced by individuals those that are worth preserving. The domain is a repository of knowledge held by the culture that preserves ideas or forms selected by the field.

In a typical cycle, an individual takes some information provided by the culture and transforms it, if the transformation is deemed valuable by society, it will be included in the domain of knowledge held by the culture, thus providing a new starting point for the next cycle of transformation and evaluation. In Csikszentmihalyi's view, creativity is not to be found in any one of these elements, but in the interactions between them.

2.1.2.Liu's Dual Generate-and-Test Model of Creativity

Recognising the need for a unified model of creativity in design computing, Liu (2000) presented a synthesis of the personal and socio-cultural views of creativity in a single model. Liu realised that the existing models of personal creativity complemented the socio-cultural models by providing details about the inner workings of the creative individual missing from the models of the larger creative system.

Liu proposed a dual generate-and-test model of creativity as a synthesis of Simon et al's model of creative thinking and Csikszentmihalyi's systems view. As its name suggests, the dual generate-and-test model of creativity encapsulates two generate-and-test loops: one at the level of the individual and the other at the level of society. The generate-and-test loop at the individual level, illustrated in Figure 2(a), provides a model of creative thinking, incorporating problem finding, solution generation and creativity evaluation. The socio-cultural generate-and-test loop, illustrated in Figure 2(b), models the interactions among the elements of Csikszentmihalyi's systems view of creativity; in particular, it captures the role that the field plays as a socio-cultural creativity test. The combined dual generate-and-test model of creativity is illustrated in Figure 2(c).

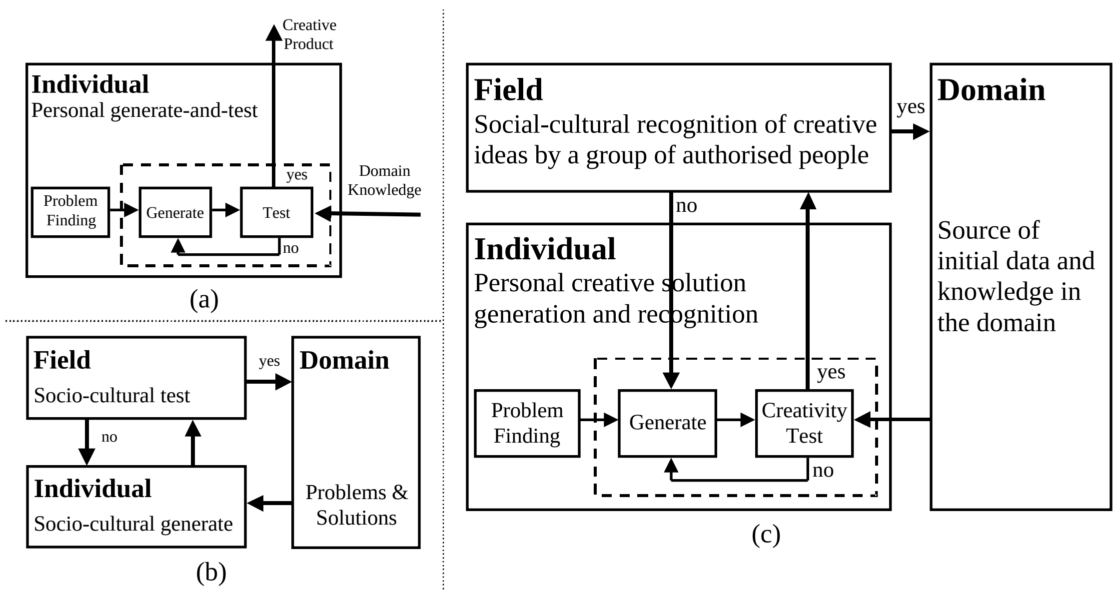

Figure 2: Liu's Dual Generate-and-Test Model of Creative Design: (a) the personal generate-and-test model, (b) the socio-cultural generate-and-test model, (c) the combined dual generate-and-test model.

Liu's model unifies Simon et al's and Csikszentmihalyi's models of creativity to form a computational model of creativity that shows how personal and socio-cultural views of creativity can be modelled in a single system. Compared to Boden's model of creativity, the dual generate-and-test model of creativity models both the P-creativity and H-creativity of individuals using the generate-and-test loops at different levels. Using the language of Gardner we may say that what distinguishes small-c creativity from big-c creativity is that big-c creativity affects changes to the domain whereas small-c creativity does not.

Liu's dual generate-and-test model shows that it is possible to cast Csikszentmihalyi's systems model in computational terms and thereby provides us with a useful basis for a framework for developing models of artificial creativity.

Before developing Liu's model further, we will examine some requirements of a computational model of artificial creativity.

3. ARTIFICIAL CREATIVITY

The artificial creativity approach that we propose here is based on Langton's approach to developing computational models of Artificial Life (Langton, 1989). The essential requirements of a computational model of Artificial Creativity are:

- The model contains a society of agents situated in a cultural environment.
- There is no agent that can direct the behaviour of all of the other agents.
- There are no rules in the agents or the environment that dictate global behaviour.
- Agents interact with other agents to exchange artefacts and evaluations.
- Agents interact with the environment to access cultural symbols.
- Agents evaluate the creativity of artefacts and other agents.

Many of the requirements of a computational model of artificial creativity are similar to the requirements of a computational model of Artificial Life. Although some of the details are different, both types of models consist of a population of agents, and both require that there are no rules or agents that can dictate global behaviour.

The additional requirement of artificial creativity not found in the requirements of Artificial Life is that the agents in an artificial creativity model must be able to make independent evaluative judgements about the creativity of agents and products in order to implement the personal and socio-cultural creativity tests found in Liu's model.

To illustrate the approach, consider how one would model a society of artists. First, we would define a repertoire of behaviours for different artistic agents and create lots of these agents. We would then start a simulation run by specifying some initial social configuration of the agents within a simulated cultural environment. From this point onwards the behaviour of the system would depend entirely on the interactions between different agents and the interactions between the agents and their cultural environment. Importantly, there would be no single agent that could enforce a definition of creativity by controlling the behaviour of all of the other agents. In addition, there would be no rules in the agents or in the environment that would define a global definition of creativity. The notions of whom and what are creative held by the society would emerge from the multiple notions of creativity held by the individual agents.

3.1. The Importance of Emergence

The requirements of artificial creativity have been designed to model the emergence of phenomena in societies of agents consistent with creativity in human society. Emergence is an important feature of artificial creativity systems, where phenomena at a certain level arise from interactions at lower levels.

In physical systems, temperature and pressure are examples of emergent phenomena. Temperature and pressure are emergent properties of large ensembles of molecules and are due to interactions at the molecular level. An individual molecule possesses neither temperature nor pressure; they are properties that only emerge when many molecules are brought together. In Artificial Life, the stable patterns in cellular automata, and the flocking behaviour of simulated birds are examples of emergent phenomena.

In artificial creativity, the socio-cultural evaluations of whom and what are creative are emergent phenomena; no individual can dictate the collective evaluations of whom and what are creative, they can only try to influence other individuals by exposing them to their products and their personal evaluations. The emergence of macro-level creativity from the interactions of individuals at the micro-level is illustrated in Figure 3.

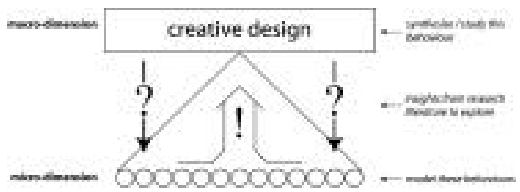

Figure 3: A behaviour-based approach to the study emergence of creative behaviour at the level of society by modelling the behaviour of individuals (after Langton, 1989).

In Boden's terms we might be tempted to say that H-creativity is emergent whereas P-creativity is not because the processes that implement P-creativity test are fixed. However, in the artificial creativity system described later the interaction between agents and the continual learning of the agents through exposure to new artefacts mean that what an agent considers to be P-creative is an emergent property of the whole system. An individual embedded within an artificial creativity system is affected by its socio-cultural context such that it will not produce the same P-creative products as it would in isolation. Hence, both H-creativity and P-creativity must be considered emergent properties of creative systems.

4. A FRAMEWORK FOR ARTIFICIAL CREATIVITY

This section presents a framework for developing computational models of artificial creativity. The framework is presented by adapting Liu's dual generate-and-test model to meet the requirements of artificial creativity listed above.

4.1. Adapting Liu's Model to Artificial Creativity

A critical aspect of Liu's model that must be addressed to develop computational models of artificial creativity is the definition of the socio-cultural creativity test. A literal implementation of Liu's model would produce a separate process that would model the socio-cultural creativity test. This is a viable solution for modelling some aspects of creativity, as demonstrated by the computational model developed by Gabora to study the memetic spread of innovations through a simulated culture Gabora (1997). Colton (2000) applied a similar socio-cultural creativity test to assess the increase in creativity due to the co-operation of agents searching a space of mathematical possibilities using different search heuristics. However, implementing a single function, or agent, that model a socio-cultural creativity test would violate one of the requirements for artificial creativity outlined previously, i.e. that no rule or agent should direct global behaviour.

Liu does not go in to details about the definition of this function but it appears that he considers this function to be outside the scope of computational models and something that can only be implemented by some form of interaction with human society. Many computational models developed reinforce this view by concentrating on the constrained generation of novel ideas in their computational models and relying on users to evaluate the creative worth of ideas. For example, see Clancey (1997) for a discussion of the social situatedness of Harold Cohen's AARON.

To computationally model the behaviour of creative societies, it is necessary to define a socio-cultural creativity test without violating the requirements of artificial creativity. The key to solving this problem is to realise that the personal creativity test inside each individual can be used to develop a socio-cultural test for creativity. The socio-cultural creativity test can be modelled by permitting the communication of artefacts and evaluations of personal creativity between individuals. An example of two individuals communicating creativity evaluations is illustrated in Figure 4.

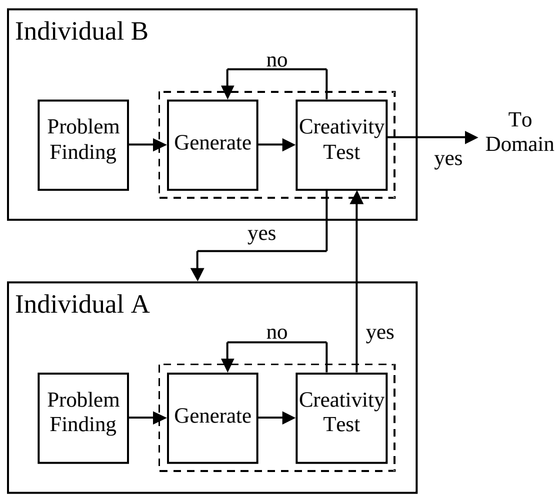

Figure 4: The communication of evaluations between individuals and its integration into the individual generate-and-test cycle.

In the interaction illustrated in Figure 4, Agent A communicates an artefact that it considers to be creative, i.e. that passes it personal creativity test, to Agent B. Agent B evaluates the artefact according to its own personal creativity test and sends its evaluation back to Agent A. In this way, Agent B can affect the generation of future artefacts by Agent A by rewarding Agent A when it generates artefacts that Agent B considers to be creative. More subtly, Agent A can affect the personal creativity test of Agent B by exposing it to artefacts that Agent A considers to be creative, because the evaluation of creativity involves an evaluation of novelty, Agent A affects a change in Agent B's notion of creativity by reducing the novelty of the type of artefacts that it communicates. By exposing Agent B to artefacts that Agent A considers to be creative, because they are novel and yet understandable, it can alter the evaluation of creativity made by Agent B.

Agent-centric evaluations of creativity permit the emergence of socio-cultural definitions of creativity as the collective function of many individual evaluations. Without agent-centric evaluations of interestingness the collection of agents would simply represent parallel searches of the same design space. To implement the socio-cultural creativity test as a collective function of individual creativity tests a communication policy is needed. A simple communication policy would be for agents to communicate a product when their evaluation of that product is greater than some fixed threshold. More complex communication policies might incorporate more strategic knowledge about when to communicate and who to communicate with.

To complete the implementation of the field as a collection of individuals, the individuals must be given the ability to interact with the domain according to some domain interaction policy. A simple domain interaction policy would follow the communication policy above and allow agents to add products of the generative process if the personal creativity evaluation is greater than a domain interaction threshold. This approach is illustrated in Figure 4. However, to ensure some level of social agreement before the addition of products to the domain, a slightly more complex domain interaction policy ensures that no individual is allowed to submit their own work to the domain. Thus, at least one other agent must find an individual's work creative before it is entered into the domain.

Making these amendments to Liu's dual generate-and-test results in the model of socio-cultural creativity illustrated in Figure 5.

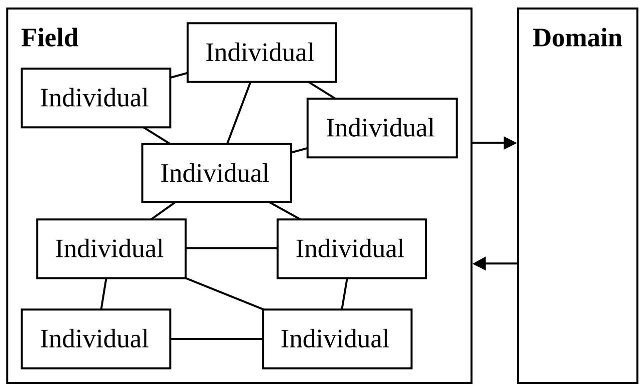

Figure 5: The Artificial Creativity model of socio-cultural creativity.

5. THE DIGITAL CLOCKWORK MUSE PROJECT

In "The Clockwork Muse" Martindale (1990) presented an extensive investigation into the role that an individual's search for novelty plays in literature, music, visual arts and architecture. He concluded that the search for novelty exerts a significant force on the development of styles.

Martindale illustrated the influence of the search for novelty by individuals in a thought experiment where he introduced "The Law of Novelty". The Law of Novelty forbids the repetition of word or deed and punishes offenders by ostracising them. Martindale argued that The Law of Novelty was merely a magnification of the reality in creative fields.

Some of the consequences of the search for novelty are that individuals that do not innovate appropriately will be ignored in the long run and that the complexity of any one style will increase over time to support the increasing need for novelty. In this section, we present a computational model of the Law of Novelty developed using our Artificial Creativity approach using curious design agents that search for novelty (Saunders & Gero, 2001b).

Our model consists of multiple "curious design agents" within a single field conducting searches for interesting and potentially creative "genetic artworks". Each agent is equipped with an evolutionary art system to allow it to generate genetic artworks and can communicate with one other agent, chosen at random, on each time step. Individuals that produce artworks that are considered creative by other agents are rewarded with "creativity credit".

5.1. The Individual: A Curious Design Agent

This subsection describes the important components of a curious design agent and the interactive evolutionary system that it interacts with. The agents in the Digital Clockwork Muse Project have been developed using a model of curiosity that we have applied to several domains can be found elsewhere (Gero and Saunders, 2000; Saunders and Gero, 2001a; 2001b; 2001c). The model of curiosity provides the essential ability for agents to evaluate the creativity of artefacts and take appropriate action, i.e. evolve new artefacts, communicate with other individuals in the field, or add an artefact to the domain.

5.1.1. Interactive Evolution

Every agent in The Digital Clockwork Muse uses an "interactive" evolutionary art system, similar to the ones devised by Dawkins, Sims, Todd and Latham, and others (Dawkins, 1987; Sims, 1991; Todd and Latham, 1992) to generate "genetic artworks". Interactive evolutionary art systems use a standard evolutionary system, e.g. a genetic algorithm, to evolve small populations of artworks that are presented to a human user for evaluation. In our system, agents take the place of human users and interact with the evolutionary art systems to search for novel genetic artworks. The flow of information between an agent and its evolutionary art system is illustrated in Figure 6.

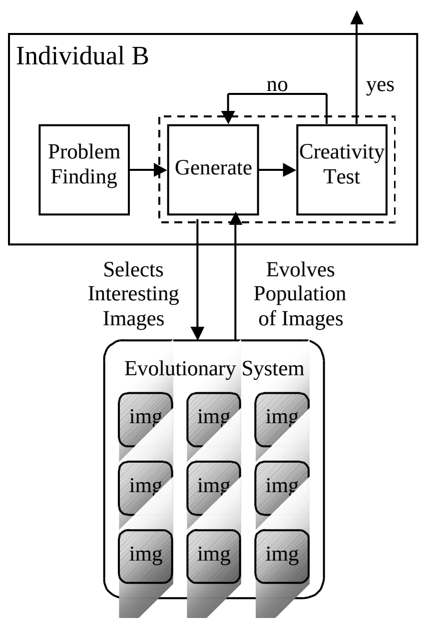

Figure 6: A curious design agent and an interactive evolutionary art system.

5.1.2. Genetic Artworks

Karl Sims is best known for his work developing one of the first interactive evolutionary art systems for complex two-dimensional bitmap images (Sims, 1991). Using a process similar to Genetic Programming, Sims devised an evolutionary art system that produced artworks by evolving symbolic function trees.

An example genetic artwork of the type evolved by the agents in this project is shown in Figure 3. This genetic artwork was evolved over the Internet as part of the International Interactive Genetic Art (IIGA) project (Witbrock and Reilly, 1999). The

IIGA project.

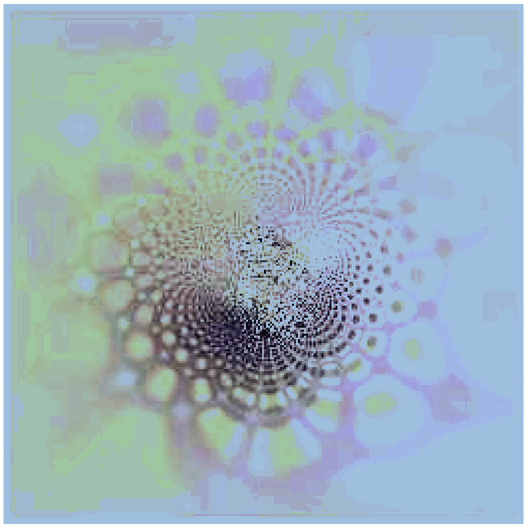

Figure 3: An example of a genetic artwork interactively evolved by a human user. (From the archive of evolved genetic artworks in Interactive Genetic Art III.)

5.1.3.Sensing

A 32x32-pixel image of each genetic artwork is analysed by a curious design agent to determine its novelty. Although this is a low-resolution image it is still large enough to allow complex artworks to be evolved. To sense the image, a relatively simple combination of a Laplacian edge-detector and a fixed intensity threshold function were used to transform a genetic artwork into a binary image, as shown in Figure 7.

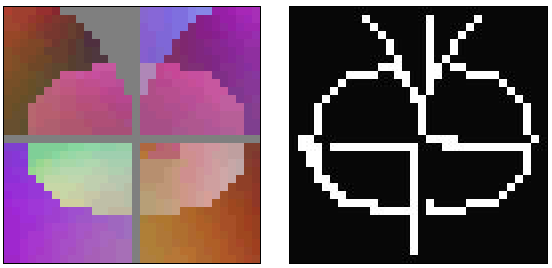

(a) (b)

Figure 7: The image processing applied to genetic artworks to extract the edge structure of the images, (a) the original image, and (b) the binary image produced by the image processing to find the most prominent edges.

5.1.4.Novelty

Each agent is equipped with a neural network to learn the categories of images as it explores the space of possible genetic artworks. A self-organising map, or SOM, (Kohonen, 1995) is used to categorise each artwork that an agent encounters into a category represented by one of the network's neurons. At each presentation of an artwork the processed binary image is converted into a vector consisting of 1024 values. As an agent explores the space of possibilities it learns a map of typical artworks for the region of the genetic art space it currently occupies. By comparing new artworks against this map, the agent can detect novel, and potentially interesting, artworks.

The map that the neural network produces provides a form of short-term memory for the agent to compare new artworks with previously created ones. The larger the network, the more neurons the agent has, and the more categories of artworks it can remember and recall for comparison.

Figure 8 shows the neighbourhoods that have formed for similar input patterns, e.g. around E2 and A6, as well as the mixing of these patterns in the intermediate areas, e.g. around D4. The mixing of representations in this way provides an agent with the ability to generalise from past experiences and hence predict aspects of unseen artefacts. This is an important ability for curious design agents because it allows them to determine the novelty of new artefacts without sampling all of the design space.

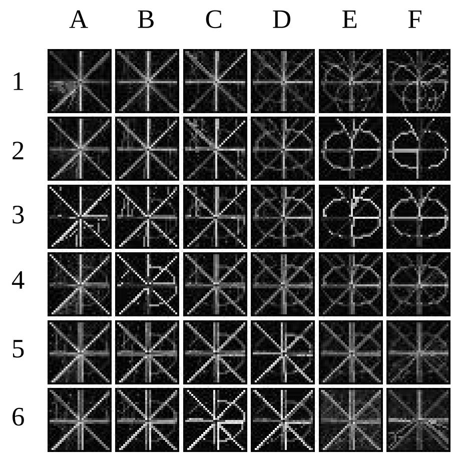

Figure 8: The prototypes represented by the 36 neurons of a self-organising map having just categorised the input shown in Figure 4b at location E2.

Novelty (N) is calculated as the categorisation error of an agent's SOM as it attempts to identify a suitable category for an artwork. Novelty values, i.e. the values of output by the best matching neuron of the neural network depend on the size of image, in this case these values are in the range N=0 and N=32, with N=0 being an exact match and N=32 being a complete mismatch. Effectively this measures the distance of the closest category prototype to the input pattern.

The Euclidean distance between the closest category prototype and a new input pattern is a rather crude measure of novelty, and more sophisticated measures have been developed by several researchers including the authors (Kohonen, 1993; Marsland et al., 2000; Saunders and Gero, 2001c), however, for the purposes of this demonstration system the measure of novelty provided by the categorisation error is sufficient and computationally inexpensive.

Novelty is used as the sole criterion to evaluate evolved artworks for interestingness. As such we define the interestingness of an artwork based on the degree to which it could not have been predicted from previous experience. This is similar to Boden's notion of P-novelty (Boden, 1990). Our definition of interestingness based on novelty alone lacks the explicit requirement for usefulness needed to model P-creativity as defined by Boden but, we argue that because interesting artworks are actionable, i.e. they promote curious action, the usefulness of an artwork is its potential to lead to other interesting artworks and is therefore, within the confines of this simple system, related to its novelty.

5.1.5. Interestingness

Interest in an artwork is calculated using an approximation to the Wundt curve, a well-known arousal response curve developed from studies of animals and humans to exposed to arousal producing stimuli, including novelty (Berlyne, 1971). The Wundt curve is sketched in Figure 9. Berlyne (1971) refers to the Wundt curve as a “hedonic function”, to indicate its relationship to the pleasure/pain response that is often associated with arousing stimuli.

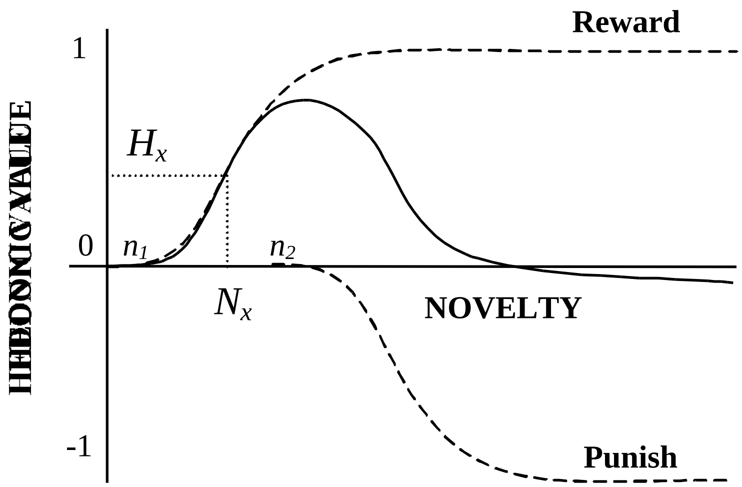

Figure 9: The hedonic function used to calculate interest. The hedonic function is shown as a solid line, the reward and punishment sigmoidal curves summed to form the hedonic function are shown dashed.

In our model the hedonic function is calculated as the sum of two sigmoidal functions whereas the Wundt curve is calculated as the sum of cumulative-Gaussian functions. The most important feature of the hedonic function used in this research that it shares in common with the Wundt curve is that it is the sum of two non-linear functions. In either case the functions are summed to produce an inverted ‘U’ shaped curve, as sketched in Figure 9. The sigmoidal function labelled ‘Reward’ represents the intrinsic reward given to the agent for finding an arousal-inducing stimulus over a fairly low threshold, n_1. The second function, labelled ‘Punish’, is the amount of punishment that the agent receives for finding an arousal-inducing stimulus over a higher threshold, n_2. By altering the thresholds for the reward and punishment sigmoid curves this peak can be positioned anywhere along the novelty axis.

The agents in the Digital Clockwork Muse use the above hedonic function to calculate the level of interest that they have in a particular artwork based upon the novelty detected by the self-organising map. Figure 9 illustrates the use of the hedonic curve with an example novelty value N_x that is mapped to its corresponding hedonic value H_x.

5.1.6. Curiosity

Through a combination of the neural network and the hedonic function the agents display a form of “curious” behaviour. Given a set of new artworks an agent will favour those that are imperfectly represented by the self-organising map, indicating the need for some learning, but are not so novel as to fall beyond the peak of the hedonic function. Thus the agent is motivated to choose artworks it has a good chance of improving its representation of by favouring similar-yet-different artworks at each time step (Berlyne, 1971). In other words, the agent shows little interest in artworks that are either too similar or too different to its previous experiences (Schmidhuber, 1991)

An agent's interest in an artwork determines the artwork's actionability. If an artwork is the most interesting at a given moment without being interesting enough to be considered creative then the artwork is selected as the starting point for further search but not sent to any other agents.

5.2. The Field: A Community of Interest

To define a field in an artificial creativity system we need to define the communication mechanisms and policies used by agents to exchange artworks and evaluations. For this project we have chosen to use the simplest implementations possible.

5.2.1.Communication

If the interestingness of an artwork breaches a threshold value that marks the lower bound of the range of potentially creative artworks then the artwork is sent to other agents for peer review.

Artworks are exchanged as messages that encode the symbolic descriptions of the artworks. Receiving agents must then express the genetic representation to recover the artwork and then evaluate it. Having expressed a received artwork an agent evaluates it according to its personal creativity test based on its own experiences. The experiences of a receiving agent are likely to be different than those of the sender and this can lead to very different evaluations of the same artwork. An artwork that was interesting for its creator may be boring to a second agent because it is too familiar or uninteresting to a third because it is not familiar enough.

An agent may find a received artwork more interesting than its own current artworks, in which case it can use the received artwork as the starting point for a new search of the genetic art space.

An advantage of passing the genetic representations of artworks between agents, rather than the artworks themselves, is that if a receiving agent finds an artwork interesting it can use the genetic representation to evolve new artworks without having to "reverse engineer" an artwork first. This is a computationally efficient approach to distributing artworks but it removes the possibility of memetic evolution of artworks through the introduction of errors during the imitation process (Dawkins, 1976). To safeguard against plagiarism and thereby stop a popular artwork being copied by all members of a population unaltered, an agent is not allowed to pass on a received artwork as its own in the same cycle; it must perform at least one evolutionary generation first.

Before using an artwork received from elsewhere an agent must pay the creator of the interesting artwork some credit, proportional to the interest the receiving agent has in the artwork. The amount of credit accumulated throughout a lifetime is used to assess how creative a particular individual is.

5.3. The Domain: A Repository for Creative Artworks

A domain is maintained by the collective actions of agents in its associated field. We have implemented the minimal domain interaction policy that ensures some form of social agreement within a field before an artwork can be added to a domain. Agents cannot add their own artworks to the domain; they can only add artworks that they receive from others. To qualify for addition to the domain an artwork must be of particularly high interestingness for the receiving agent, most likely higher than that required for an artwork to be considered worthy of communicating to another member of is field. If so, the artwork is added to the domain with a label indicating the agent that created it.

Future generations of genetic artists begin their search with artworks that have been added to the domain, however, the dynamic nature of the socio-cultural evaluation process means that artworks that were considered creative are likely to be no longer considered creative because they are too familiar to the field. Therefore, the domain does not provide instant access to creative works, but rather a store of familiar starting points from which new creative artworks can be produced. The real advantage of starting with artworks stored in the domain is that they are already familiar to other members of the field. The result of a short search for novel artworks starting with examples from the domain is likely to be new artworks that are similar-yet-different with respect to the domain, making them ideal candidates for being creative.

Researchers of artificial creativity can also use the records kept in the domain as a means to trace the development of artistic styles considered creative over time.

6. EXPERIMENTS IN ARTIFICIAL CREATIVITY

The following experiments were conducted with the aim of confirming Martindale's predictions for Artificial Creative systems and to investigate other interesting emergent behaviour.

6.1. The Law of Novelty

We investigated the effects of the search for novelty, by producing agents with different hedonic functions. The aim was to show that agents are not recognised as creative when they fail to innovate inappropriately. Agents can innovate inappropriately either by producing "boring" images that are too similar to images previously experienced by other agents, or by producing "radical" images that are too different for other agents to appreciate.

We have simulated both types of inappropriate innovation in a single simulation. For this experiment we created a group of agents most of whom, agents 0-9, shared the same hedonic function, i.e. the same preference for average novelty (N=11). Two of the agents have quite different novelty preferences. One, agent 10, has a preference for low amounts of novelty (N=3) and the other, agent 11, has a preference for high amounts of novelty (N=19). Agents with a lower novelty preference tend to innovate at a slower rate than those with a higher hedonic preference. The results of the simulation are presented in Table 1.

Table 1: The attributed creativity for a group of agents with different preferences for novelty.

Agent ID	*Preferred Novelty*	*Attributed Creativity*
0	N=11	5.43
1	N=11	4.49
2	N=11	4.50
3	N=11	3.60
4	N=11	4.48
5	N=11	1.82
6	N=11	6.32
7	N=11	8.93
8	N=11	10.72
9	N=11	5.39
10	N=3	0.0
11	N=19	0.0

The results show the agents with the same preference for novelty to be somewhat creative according to their peers, with an average attributed creativity of 5.57. However, neither agent 10 nor agent 11 received any credit for their artworks. Consequently none of the artworks produced by these agents were saved in the domain for future generations. When these agents expired nothing remained in the system of their efforts.

The results show that while an agent must innovate to be considered creative, it must do so at a pace that matches other agents to achieve recognition. The agent with a preference for high levels of novelty and hence rapid innovation was just as unsuccessful in gaining recognition as the agent with a low novelty threshold that innovated too slowly.

6.2. The Emergence of Cliques

We have also investigated the behaviour of groups of agents with different hedonic functions. To do this we created a group of 10 agents, half of them had a hedonic function that favoured novelty N=6 and the other five agents favoured novelty values close to N=15. Figure 9 shows the payments of creativity credit between the agents in recognition of interesting artworks sent by the agents.

Receiver \ Sender	0	1	2	3	4	5	6	7	8	9
0		2	8	1	2					
1										
2	2	1		1	3					
3	4	5	2		5					
4	2	3	3	2						
5					1		6	1	3	5
6						3		4	5	1
7						3	5		1	4
8						4	3	2		4
9						1	4	4	4	

Figure 9: A matrix showing the total number of messages carrying credit for being creative between the agents of the simulation.

Two areas of frequent communication can be seen in the matrix of payment messages shown in Figure 9. The agents with the same hedonic function frequently send credit for interesting artworks amongst themselves but rarely send them to agents with a different hedonic function. There are a large number of credit messages between agents 0-4 and agents 5-9, but only one payment between the two groups – agent 4 credits agent 5 for a single interesting artwork.

The result of putting collections of agents with different hedonic functions in the same group appears to be the formation of cliques: groups of agents that communicate credit frequently amongst themselves but rarely acknowledge the creativity of agents outside the clique. As a consequence of the lack of communication between the groups the style of artworks produced by the two cliques also remains distinct.

Communication between cliques is rare but it is an important aspect of creative social behaviour. Communication between cliques occurs when two individuals in the different cliques explore design subspaces that are perceptually similar. Each of the individuals is then able to appreciate the other’s work because they have constructed appropriate perceptual categories. The transfer of artworks from a source to a destination clique will introduce new variables into the creative processes of the destination clique, the two cliques can then explore in different directions, just as two individuals do when they share artworks. Cliques can therefore act as “super-artists”, exploring a design space as a collective and communicating interesting artworks between cliques.

Figure 10 is a screenshot of the running simulation that has formed two cliques. To help visualise the emergent cliques, the distances between agents are shortened for agents that communicate frequently. The different styles of the two groups can also be seen, with agents 0-4 producing smooth radial images with low a fractal dimension (~1.4) and agents 5-9 producing fractured images with clearly defined edges and a higher fractal dimension (~1.7).

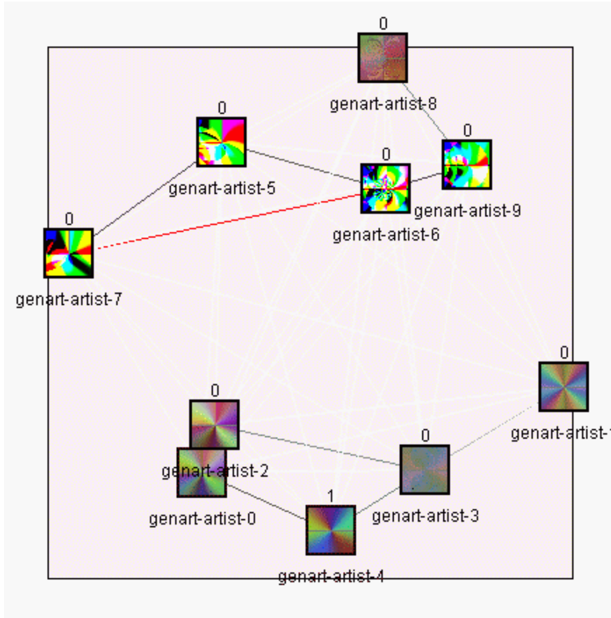

Figure 10: A screenshot of the simulation clearly showing the two cliques. The squares represent agents. The images show the currently selected genetic artwork for each agent. The number above each square shows the agent's attributed creativity. The dark lines between agents indicate the communication of credit.

The stability of these cliques depends upon how similar the individuals in different subgroups are and how often the agents in one subgroup are exposed to the artworks of another subgroup. Cliques have been observed in simulations with agents that share similar hedonic functions, but the cliques are often consist of 3-4 individuals and do not last for long before the clique splits apart and the agents form new cliques with other agents. Further research is needed to determine whether other factors of individual behaviour can similarly affect the social structure.

7. FUTURE RESEARCH

The artificial creativity framework implemented here provides several opportunities for developing future models of social creativity. Three possible directions for future work are: (1) the simulation of larger creative societies, (2) the development of new types of agents, and (3) the development of more complex social interactions.

7.1. Large Creative Societies

The ability to simulate larger creative societies will permit the study of the spread of innovations (Gabora, 1997; Goldenberg et al., 2000) and styles. It may also facilitate the emergence of new fields as cliques attain a critical size. Spatial and topological relationships will become more important issues in large population models.

7.2. New Types of Agents

There are several other important players in creativity societies besides the producers of innovations including, e.g. consumers, distributors, critics, etc. Each has their own role to play in artificially creative societies; consumers evaluate products, distributors distribute products widely, and critics distribute their evaluations widely. These roles are illustrated in Figure 11.

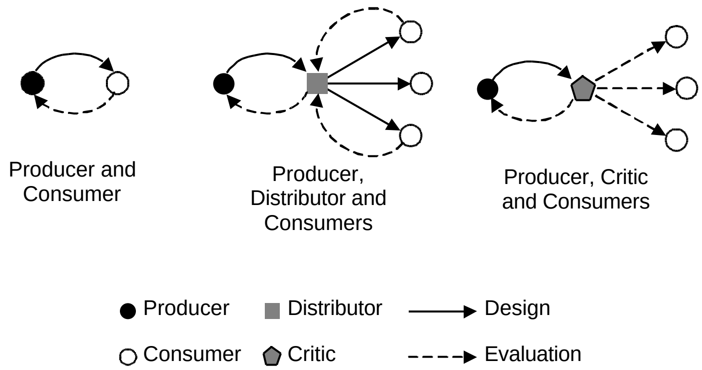

Figure 11: Three different types of individuals and their roles in the communication of designs and evaluations in creative design societies.

Convincing other people that you've had a creative idea is often harder than having the idea in the first place (Csikszentmihalyi, 1999). In non-homogenous societies of agents, the selection of which agents to communicate with becomes important for agents seeking recognition from their peers.

7.3. Strategic Knowledge

Simulations of technological innovation in industry show that the consideration of the costs of innovation in decision-making can lead to complex behaviour (Haag and Liedl, 2001). Simulating similar costs in the design process may provide a better understanding of the economics of creative design in creative societies and the strategies needed to manage creativity with limited resources.

8. CONCLUSIONS

The computational work presented in this paper has illustrated the artificial creativity approach to developing models of creative societies. By adapting Liu's dual generate-and-test model of creativity we have produced a model of creative societies that can be used to study socio-cultural creative behaviour as an emergent property arising from the creative behaviour of individuals. The implemented system models the evolution of notions of creativity within an artificial society over time as individuals come and go, the field changes in composition, and the domain is altered.

The emergence of social behaviour, e.g. The Law of Novelty, and dynamic social structures, e.g. cliques; suggest that the artificial creativity approach to developing models of creative societies may contribute new insights into the nature of creative design in socio-cultural situations. Figure 12 illustrates the different levels at which creativity may be studied as a pyramid of emergent properties. Each level represents a different aspect of creativity that is emergent from the ones below it. The foundations of the creative pyramid are the processes internal to the creative agent that allows it to generate-and-test ideas. The result of executing these processes is the creative products. Traditionally, computational research has concentrated on these two levels by encoding processes thought to be important in creativity in a piece of software and getting experts to examine the results of running those processes to determine whether the processes are creative. In traditional computational models, the higher levels of the pyramid are not modelled in the software and are provided by people.

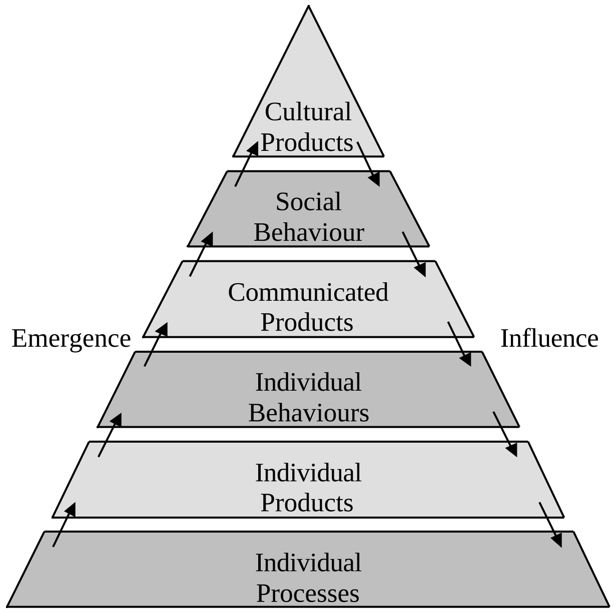

Figure 12: A pyramid of creativity.

Artificial creativity suggests a different approach; instead of evaluating the products of a piece of software to determine its creativity, it focuses upon the behaviours of agents and artificial societies. Artificial creativity is concerned with modelling the creative behaviours of individuals, e.g. curiosity, and studying the emergent social behaviours when individuals are put together. Because individuals in an artificial creativity simulation must be able to evaluate the creativity of communicated products and hence other individuals, the details of the products of individuals become less important. More important in the study of artificial creativity are the socio-cultural structures that emerge as a consequence of the communication of products and evaluations.

The artificial creativity approach presented here permits the computational study of highest levels of creativity illustrated in Figure 12 without having to develop agents that can integrate, and achieve creative status, in human society. Artificial creativity simulations permit the experimentation with creativity in artificial societies that would be impossible in the real world, allowing the study of *creativity-as-it-is* in the context of *creativity-as-it-could-be*.

ACKNOWLEDGEMENTS

This research is supported by an Overseas Postgraduate Research Scholarship and by a University of Sydney Postgraduate Award. We wish to thank Michael Witbrock, Scott Neal Reilly and John Mount for developing the inspirational evolutionary art systems International Interactive Genetic Art I, II and III. In particular we would like to thank John Mount for allowing us to use his source code to develop the evolutionary art system, without which the hedonic value of this paper would be lower. We also wish to thank Ricardo Sosa, Julie Jupp and Hsien-Hui (Michael) Tang for their insights and support developing the ideas presenting here. We would also like to thank Ricardo Sosa for allowing us to use his illustration in Figure 3.

REFERENCES

Berlyne, D. E. (1971) *Aesthetics and Psychobiology*, Appleton-Century-Crofts, New York.

Boden. M. A. (1990) *The Creative Mind: Myths and Mechanisms.* Cardinal, London.
Clancey, W. J. (1997) *Situated Cognition: On Human Knowledge and Computer Representations*, Cambridge University Press, Cambridge, England.
Colton, S., Bundy, A. and Walsh, T. (2000) Agent Based Cooperative Theory Formation in Pure Mathematics, *in* G. Wiggins (ed.), *Proceedings of AISB 2000 Symposium on Creative and Cultural Aspects and Applications of AI and Cognitive Science*, pp. 11–18, Birmingham, UK.
Csikszentmihalyi, M. (1988) Society, culture, and person: a systems view of creativity, *in* R. J. Sternberg (ed.), *The Nature of Creativity,* Cambridge University Press, Cambridge, UK, pp. 325–339.
Csikszentmihalyi, M. (1999) Implications of a Systems Perspective for the Study of Creativity, *in* R. J. Sternberg (ed.), *Handbook of Creativity*, Cambridge University Press, Cambridge, UK, pp. 313–335.
Dawkins, R. (1976) *The Selfish Gene*, Oxford University Press, Oxford.
Dawkins, R. (1987) *The Blind Watchmaker*, Norton and Co., London.
Gabora, L. (1997) The origin and evolution of culture and creativity. *Journal of Memetics: Evolutionary Models of Information Transmission* **1**(1): 1–28.
Goldenberg, J., Libai, B., Solomon, S., Jan, N. and D. Stauffer (2000) Marketing percolation, *Physica A*, **284**(1-4): 335–347.
Gardner, H. (1993) *Creating Minds*, Basic Books, New York.
Gero, J. S. and Reffat, R. (1997) Multiple representations for situated agent-based learning, *in* B. Varma and X. Yao (eds.) *ICCIMA'97*, Griffiths University, Australia pp. 81–85.
Haag, G. and Liedl, P. (2001) Modelling and simulating innovation behaviour within micro-based correlated decision processes, Journal of Artificial Societies and Social Simulation, **4**(3), http://www.soc.surrey.ac.uk/JASSS/4/3/3.html.
Hofstadter, D. (1979) *Godel, Escher, Bach: An Eternal Golden Braid*, Basic Books, New York.
Hofstadter, D. and the Fluid Analogies Research Group (1995) *Fluid Concepts & Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*, Basic Books, New York.
Koestler, A. (1964) *The Act of Creation,* Macmillan, New York.
Kohonen, T. (1993) *Self-Organization and Associative Memory*, 3rd ed., Springer, Berlin.
Kohonen, T. (1995) *Self-Organizing Maps*, Springer–Verlag, Berlin.
Langley, P., Simon, H. A., Bradshaw, G. L., and Zytkow, J. M. (1987) *Scientific Discovery: Computational Explorations of the Creative Processes,* MIT Press, Cambridge, MA.
Langton, C. G. (1989) Artificial life, *Artificial Life*, Addison–Wesley, pp. 1–47.
Liu, Y. T. (2000) Creativity or novelty? *Design Studies* **21**(3): 261–276.
Marsland, S., Nehmzow, U. and Shapiro, J. (2000) A real–time novelty detector for a mobile robot, in *Proceedings of EUREL European Advanced Robotics Systems Conference,* Salford, England.
Martindale, C. (1990) *The Clockwork Muse*, Basic Books, New York.
Newell, A., Shaw, J. C. and Simon, H. A. (1962) The process of creative thinking, *in* H. Gruber, G. Terrell and M. Wertheimer (eds.), *Contemporary Approaches to Creative Thinking*, Atherton Press, New York.
Partridge, D. Rowe, J. (1994) *Computers and Creativity*, Intellect Books, Oxford.
Saunders, R. and Gero, J. S. (2001a) Designing for Interest and Novelty: Motivating Design Agents, *in Proceedings of CAAD Futures 2001,* Eindhoven.
Saunders, R. and Gero, J. S. (2001b) A Curious Design Agent: A Computational Model of Novelty-Seeking Behaviour in Design, *in Proceedings of the Sixth Conference on Computer Aided Architectural Design Research in Asia (CAADRIA 2001)*, The University of Sydney, Australia.
Saunders, R. and Gero, J. S. (2001c) The Digital Clockwork Muse: A Computational Model of Aesthetic Evolution, *in* G. Wiggins (ed.), *Proceedings of the AISB'01 Symposium on AI and Creativity in Arts and Science*, SSAISB, York, UK.
Schmidhuber, J. (1991) A possibility for implementing curiosity and boredom in model-building neural controllers, *in* J. A. Meyer and S. W. Wilson (eds.) *Proceedings of the International Conference on Simulation of Adaptive Behaviour: From Animals to Animats*, MIT Press/Bradford Books, pp. 222–227.
Simon, H. A. (1981) *The Sciences of the Artificial*, MIT Press, Cambridge, MA.
Sims, K. (1991) Artificial evolution for computer graphics. *Computer Graphics* **25**(4): 319-328.
Taylor, C. W. (1988) Various approaches to the definition of creativity, *in* R. J. Sternberg (ed.), *The Nature of Creativity,* Cambridge University Press, Cambridge, UK, pp. 99–124.
Todd, S. and Latham, W. (1992) *Evolutionary Art and Computers*. Academic Press, London.
Witbrock, M. and Reilly, S. N. (1999) Evolving genetic art, *in* Bentley, P. J. (ed.) *Evolutionary Design by Computers*, Morgan Kaufman, San Francisco, CA, pp. 1–73.

The full reference for this paper is:

Saunders, R and Gero, JS (2001) Artificial creativity: A synthetic approach to the study of creative behaviour, *in* JS Gero and ML Maher (eds), *Computational and Cognitive Models of Creative Design V*, Key Centre of Design Computing and Cognition, University of Sydney, Sydney, pp. 113-139.