

# THE DIGITAL CLOCKWORK MUSE: A COMPUTATIONAL MODEL OF AESTHETIC EVOLUTION

Rob Saunders; John S. Gero

Key Centre of Design Computing and Cognition; University of Sydney, NSW 2006, Australia  
rob@arch.usyd.edu.au; john@arch.usyd.edu.au

## Abstract

This paper presents a computational model of creativity that attempts to capture within a social context an important aspect of the art and design process: the search for novelty. The computational model consists of multiple novelty-seeking agents that can assess the interestingness of artworks. The agents can communicate to particularly interesting artworks to others. Agents can also communicate to reward other agents for finding interesting artworks. We present the results from running experiments to investigate the effects of searching for different degrees of novelty on the artworks produced and the social organisation of the agents.

## 1 Introduction

In “The Clockwork Muse” 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 force on the development of artistic movements. Martindale illustrated the influence that the search for novelty has on creative activity with the following thought experiment, “The Law of Novelty”:

We live in a predictable world. [...] Every morning we are bombarded with automatic “Good mornings” and routine inquiries about how we are. Our days are full of pat questions and equally pat replies. Imagine what would happen, however, if some of us, tiring of this state of affairs, decide to do something about it. Decide, indeed, to outlaw any and all repetition. Once something has been said, it can never be said again. Once something has been done, it cannot be done again. This requires no act of Congress. We can implement it ourselves. We can, as well, impose a sanction more severe than the death penalty. Anyone who says something that has already been said, or does something that has already been done, will simply be ignored. Someone who persists in the crime of repetition will find that he or she ceases to exist. No one will pay the slightest attention to the person.

Martindale elaborated some of the consequences of adopting his Law of Novelty. For example, discourse would be forced into ever more concrete and specific forms to en-

sure that the meaning is understood while at the same time driving speakers to circumlocution in order to avoid repetition. For example, a salesman’s description of a toaster becomes an exercise in metaphors: “Rather than saying, ‘This is our newest model,’ something such as ‘This is the rising sun of the destroyers of leavened moisture,’ is called for.”

Martindale argued that what he described was not a revolution but merely a magnification of the world we live in and that in fact the Law of Novelty is applied in its purest form in the arts. Exact replication is not allowed in any of the arts: otherwise the notion of forgery would be a non-sense. And although artists may not be ignored while they are alive the rule of novelty takes effect with exceptionless brutality upon their deaths. It is then that they begin to be ignored. This paper presents an attempt to implement the Law of Novelty in a computational model of creativity.

## 2 Previous Work

In laboratory experiments Martindale has shown that the transmission of concepts through the imitation of drawings introduces errors that are reminiscent of the evolution of “memes” introduced by Dawkins and subsequently elaborated by Blackmore and others (Dawkins, 1976; Blackmore, 1999; Gabora, 1997). Martindale calls this process aesthetic evolution.

Gabora has developed a memetic theory of creativity that stresses the important relationship between innovation and imitation in the spread of creative ideas and cultural evolution (Gabora, 2000). Gabora has also developed a

computational model, “Memes and Variations”, that demonstrates this theory for a fixed fitness function.

In “Memes and Variations” agents exchange information about ritual dance moves by imitating the movements of nearby agents. The success of any innovations made by the agents as a result of mutations are evaluated using an objective fitness function that calculates the number of correct limb positions over time. By using a combination of simulation, imitation and innovation Gabora showed how successful innovations quickly spread through a population of agents.

Colton et al. (2000) investigated a quite different type of culturally situated creativity in a study of agent based cooperative theory formation in pure mathematics. Colton et al. compared the performance of groups of collaborating agents with single agents. They discovered that small groups of collaborating agents with different search strategies outperformed single agents under a number of criteria. As part of their investigation, Colton et al. developed a definition of creativity appropriate to theory formation, based on the novelty of a theory’s categorisation of a set of numbers. They used this measure of creativity to assess the relative performance of collaborating groups and found that larger groups with more diverse search strategies were more creative.

The computational model of creativity presented in this paper attempts to use a different approach to modelling creativity from previous models with fixed, objective definitions of creativity by supporting the emergence of socially defined notions of what and who are creative. Our model incorporates Martindale’s insights about the importance of the search for novelty into a systems view of creativity that allows the determination of what products are creative and which agents are creative to emerge from the interaction of agents. Thus we call our model “The Digital Clockwork Muse” in recognition of the constant drive to search for novelty placed upon the creative individual by the relentless need to innovate in order to achieve recognition in a social context.

In the following section we examine previous approaches to defining interestingness in relation to novelty. We then describe “The Digital Clockwork Muse” before giving some experimental results in Section 5. We conclude with a discussion of possibly future work and the potential applications of a similar model of social creativity to support artists and designers.

### 3 Interestingness

The need for a notion of interestingness was recognised early in the development of computational models of mathematical and scientific discovery (Lenat, 1976, 1983; Langley et al., 1987). More recently the development of data-mining techniques for knowledge discovery in databases has lead to a similar need (Silberschatz and Tuzhilin, 1996).

Early models of discovery used various ad hoc rules to define what was interesting and to propagate these evaluations of interestingness among related concepts. For example, of the 242 heuristics used in Lenat’s AM, a total of 43 heuristics were designed to assess the interestingness of a concepts, 33 of which were concerned with propagating values originated elsewhere. As Colton (2000) notes “AM could make a little interest go a long way”. Despite a long history of research into knowledge discovery in AI, a definition of interestingness remains elusive.

Silberschatz and Tuzhilin (1996) suggest that a definition of interestingness can be either objective or subjective: objective interestingness uses relationships found entirely within the object considered interesting, while subjective interestingness compares properties of the object with beliefs of a user to determine interest. Two aspects that make something subjectively interesting are that it is unexpected and/or actionable. Unexpectedness depends upon an agent’s ability to predict an as-yet-unseen event. Actionability depends upon whether an agent can take action as a consequence of a discovery. The two concepts are conceptually independent although Silberschatz and Tuzhilin note that unexpected discoveries are often actionable.

For the purposes of our research we use a subjective notion of interestingness to guide the exploration of a space of possible genetically produced artworks (genetic artworks). The subjective measure of interestingness uses a measure of unexpectedness, or novelty, to select which of the artworks available at an instance in time should be used to continue its search for further novelty and in its attempts to get itself recognised as creative. The possible actions available to an agent when dealing with a novel artwork are described in Subsection 5.5.

### 4 The Digital Clockwork Muse

Our model of culturally situated creativity consists of multiple agents within a single field conducting searches for interesting genetic artworks. Each agent is equipped

with an evolutionary art system to allow it to generate genetic artworks. For each artwork generated the agent must assess the novelty of the image by sensing important features of the image. The sensed image is then categorised. All colour information is discarded by the image processing involved in sensing and so the agents in this system are interested only in the structure of the images evolved. The error in categorisation of an artwork is used as the novelty of an artwork and from this measure the interest that the agent has in an artwork if calculated. The agent uses the interest calculated for each artwork to determine what actions to take in terms of selecting artworks to evolve and selecting artworks to exchange with other agents. The flow of information between an agent and its evolutionary art system is illustrated in Figure 1.




Figure 1: An agent and an “interactive” evolutionary art system showing the signals sent between the agent and the evolutionary system.


When an individual finds an interesting genetic artwork it may choose to present the artwork to other agents for peer review if the artwork is thought to be interesting enough. If another agent agrees that an artwork is interesting it can choose to use it as a starting point for its own search for novelty. If it decides to do so it must first credit the original creator with some creativity in finding an interesting artwork. The crediting of the creator does not change its abilities but it does permit the interesting artwork to be added to store of interesting examples with the creator’s name attached for future individuals to use as a starting point for their searches.

The following subsections describe the major components and behaviours of the agents developed and their affects on the collective definition of creativity. We begin with a brief review of evolutionary art systems. The internal processes of sensing, learning, novelty detection and the calculation of interestingness are covered in Subsections 4.2-4.4. Subsection 4.5 concentrates on the behaviour of agents as a consequence of the discovery of interesting artworks.

## 4.1 Genetic Art

Every agent in The Digital Clockwork Muse has an “interactive” evolutionary art system, similar to the ones devised by Dawkins, Sims, Todd and Latham, and others (Dawkins, 1986; Sims, 1991; Todd and Latham, 1992). In these systems the agents take the place usually held by human users and interact with the evolutionary art systems to search for novel genetic artworks.

Interactive evolutionary art systems work by using a standard evolutionary system, e.g. a genetic algorithm, to evolve small populations of artworks that are presented to the user for evaluation. The user of the evolutionary art system can then evaluate the artworks based on some form of aesthetic preference and instruct the evolutionary art system to generate a new population of artworks based on those preferences.



$$(+ (\sin x) (* (\text{rand } x y) x))$$

Figure 2: An example function tree and its corresponding Lisp expression for a genetic artwork.

Dawkins first popularised this method of evolving aesthetically pleasing images in his book “The Blind Watchmaker” (Dawkins, 1986) with a program that evolved “biomorphs” – small stick figures that resembled insects, butterflies or trees depending on the specifics of the evolved genes<sup>1</sup>. Todd and Latham (1992) developed a much more complex form of interactive evolution to allow William Latham to evolve “virtual sculptures” that he has exhibited in galleries worldwide. Karl Sims is probably 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 evol-

<sup>1</sup> See Dawkins (1989) for an interesting discussion on the evolution of evolvability detailing the development of the biomorph program.

ing symbolic function trees. A simple function tree is illustrated in Figure 2.

To produce an image a symbolic function tree is evaluated at a set of points, typically between  $(-0.5, -0.5)$  and  $(0.5, 0.5)$ , that corresponds to every pixel location in an image of a given size. The values for  $x$  and  $y$  at the terminal nodes of the tree are substituted with the  $x$  and  $y$  co-ordinates of the sample point and the values for every higher node in the tree are recalculated. The root node evaluation is then interpreted to produce a colour value that is assigned to the pixels of the output image.

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



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.)

Unlike the work of Sims, which uses a rich mix of computer graphics procedures and image processing techniques, the genetic art systems of the IIGA project use quaternion mathematics that deal with four-dimensional numbers. The result of evaluating a quaternion expression, i.e. evaluating the quaternion function tree at a four-dimensional location corresponding to a pixel location on the image plane, is a four-dimensional number that can be transformed into a three-dimensional vector describing the red, green, and blue components of a pixel colour. The agents in The Digital Clockwork Muse use the same code as that used in the IIGA projects to evolve images over the Internet.

## 4.2 Image Processing

A  $32 \times 32$ -pixel image of each genetic artwork is produced for analysis and categorisation in order to determine its novelty. Although this is a low-resolution image it is still large enough to allow complex artworks to be evolved.

To be able to assess the novelty of a genetic artwork some aspects of it must first be sensed using image-processing techniques. The choice of image processing technique depends on what are the most important aspects of the image for categorisation purposes. A relatively simple combination of a Laplacian edge-detector and a fixed intensity threshold function were used to extract a binary image of the predominant edges in an artwork, as shown in Figure 4b.




Figure 4: 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.

## 4.3 Learning and Novelty

Each agent is equipped with a neural network and is capable of learning as it explores the space of possible genetic artworks. Consequently, 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 novel, and potentially interesting, artworks are detected.

A neural network called a self-organising map, or SOM, (Kohonen, 1993) 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. This vector is compared with an equal length vector for each neuron in the SOM, often referred to as the weights of the neuron. A Euclidean distance between the input vector and the neuron weights is calculated and the neuron that has the closest vector of

weights is declared the winner. The winning neurons weights are then updated to reduce the distance between them and the input vector. In addition nearby neighbours are also updated to reduce the distance between their weight vectors and the input vectors. In our agents we use a Gaussian neighbourhood function to reduce the amount that the weights of neighbouring neurons are updated in proportion to their distance to the winning neuron.

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. The neurons of the SOMs used are arranged in a two-dimensional square lattice and range in size from 36 neurons to 100 neurons. Figure 5 shows the neighbourhoods that have formed for similar input patterns, e.g. around E2 and A5, as well as the mixing of these patterns in the intermediate areas, e.g. around D4.




Figure 5: 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 a newly encountered artwork. Effectively this measures the distance of the closest category prototype to the input pattern. This 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, to appear). However, for the purposes of this system the measure of novelty provided by the categorisation error is sufficient and computationally inexpensive.

The reported novelty values in the remainder of this paper are the raw novelty values, i.e. the values of output by the best matching neuron of the neural network. For the size of image used these values range between  $N=0$  and  $N=32$ , with  $N=0$  being an exact match and  $N=32$  being a complete mismatch.

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 novelty being more than simple newness as part of her definition of creativity (Boden, 1990). However, Boden also requires that the products of creativity be useful. Our definition of interestingness based on novelty alone lacks this explicit requirement for usefulness, although one could argue that, because interesting artworks are actionable, the usefulness of a genetic artwork for the agents in this system is its potential to lead to other interesting artworks.

#### 4.4 Interestingness

Interest in an artwork is calculated using an approximation to a well-known arousal response curve from studies of animals and humans to various forms of arousal. The arousal response curve is called the Wundt curve. An approximation to the Wundt curve is sketched in Figure 6 and described in some detail in Berlyne (1971). Berlyne also refers to the Wundt curve as a "hedonic function", in reference to the pleasure/pain response that is often associated with responses to arousing stimuli.




Figure 6: 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.

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 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. In either case the functions are summed to produce an inverted ‘U’ shaped curve, as sketched in Figure 6.

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 is a punishment for finding an arousal-inducing stimulus over a higher threshold,  $n_2$ .

The agents 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 6 illustrates the use of the hedonic curve with an example novelty value  $N_x$  that is used to calculate its corresponding hedonic value  $H_x$ . The preferred degree of novelty for an agent is determined by the position of the peak on the hedonic curve along the novelty axis. By altering the thresholds for the reward and punishment sigmoid curves this peak can be positioned anywhere along the novelty axis.

#### 4.5 Curiosity and Actionability

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

However, if the agent is so interested in an artwork as to breach a threshold value that marks the lower bound of the range of potentially creative artworks then the artwork will be sent to other agents for peer review.

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)

Upon receiving an artwork an agent evaluates it according to its own experience that will most likely differ from that of the originator. 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. Alternatively, an agent may find a received artwork more interesting than its own current offerings, in which case it will use the received artwork as the starting point for a new search of the genetic art space for interesting new artworks.

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.

It is also incumbent upon the receiving agent to add a record of the interesting artwork and the creating artist to a store of creative examples for posterity. Future generations of genetic artists can thus begin their search with artworks that were once considered creative, although they may no longer be depending on the intervening developments in the field. The record of interesting artworks can be used as a means to trace the development of artistic styles considered creative over time.

## 5 Experimental Results

In this section we report on some experiments we have conducted using agents in The Digital Clockwork Muse environment. First we report on experiments conducted to examine the relationship between the position of the hedonic function described above along the novelty axis and the complexity of the images produced by the agents using the function.

### 5.1 Experiment 1: Hedonic Complexity

Our first experiment investigated the relationship between the hedonic function that drives the process of evolving genetic artworks and the complexity of the artworks produced. To measure the complexity of the images we calculated the fractal dimension of the images used to train the neural networks, i.e. after image processing to find the dominant edges. The fractal dimension was estimated using the box counting method.

To investigate the relationship between the preferred degree of novelty and the fractal dimension of the resulting images we created two separate groups of agents, one group that preferred novelty values of  $N=18$  and another that favoured novelty values of  $N=11$ . Both groups were allowed to run for 50 time steps in total.

Figure 7 shows how the average fractal dimension of the selected images by the three agents in each group changed

over time. The graph clearly shows that agents with a preference for greater novelty tend towards producing images with higher fractal dimensions.




Figure 7: The development of two distinct styles of images with different fractal dimensions in two groups of agents with hedonic functions that peak for the values of novelty indicated.

To confirm this relationship between the fractal dimension of the images produced and the preference for novelty the same test (3 agents/group for 50 time steps) was performed for a range of novelty values. Figure 8 shows that the relationship between the preferred value of novelty and the fractal dimension of the resulting images is almost linear for the large proportion of values for preferred novelty.




Figure 8: A comparison of the average values for 3 agents of the fractal dimension of evolved images after 50 time steps against a range of peak hedonic values. The equation and dashed line show the result of performing a linear regression on the sample points.

Performing a linear regression on the data points we discover that on average the fractal dimension of the resulting image goes up by 0.1 per unit step in novelty preferred. Visually this means that the images produced by agents that prefer greater novelty appear more complicated than those produced by agents that prefer lower amounts of novelty, see Figure 11 for some examples of artworks evolved by agents with different hedonic functions.

How can we explain this relationship between the preferred novelty of an agent and the fractal dimension of the resulting images? One explanation is that the curious exploration of the space of genetic artworks drives the agents towards subspaces that have an appropriate amount of local variability to continually satisfy the need for novelty. Consequently, agents that prefer novel forms will tend towards areas of the design space that produce more complex images, as there is a great deal more variability in complex images than in simple ones.

The results of this experiment confirm Martindale’s hypothesis for our groups of agents: the search for greater novelty produces more complex forms.

## 5.2 Experiment 2: The Law of Novelty

In our second experiment we investigated the effects of applying The Law of Novelty on agents with different hedonic functions. The aim of this experiment was to show that agents are isolated when they fail to innovate in ways that other agents can appreciate.

The failure to innovate appropriately may be because the agent produces “boring” images that are too simple, and hence have already been learned by other agents. Alternatively, an agent may fail to innovate appropriately because it artworks that are too radical and that no other agent can understand.

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

We have simulated both types of inappropriate innovation in a single simulation of The Law of Novelty. 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). The results of the simulation are presented in Table 1.

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 with that preferred N=3 nor agent 11 that preferred N=19 received any credit for their artworks. Consequently none of the artworks produced by these agents were saved in the store of example artworks 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 the other agents that it must communicate with in order 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 did not innovate.

### 5.3 Experiment 3: Novelty Cliques

Our final experiment investigated the behaviour of groups of agents with incompatible hedonic functions that are placed in the same social setting. 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.

k. Receiver	Sender									
	0	1	2	3	4	5	6	7	8	9
0		2	8	1	2					
1										
2	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 show 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 incompatible 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.

Figure 10 is a screenshot of the running simulation that clearly shows the two cliques formed. The distances between agents are shortened for agents that communicate frequently. The different styles of the two groups can be clearly 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).

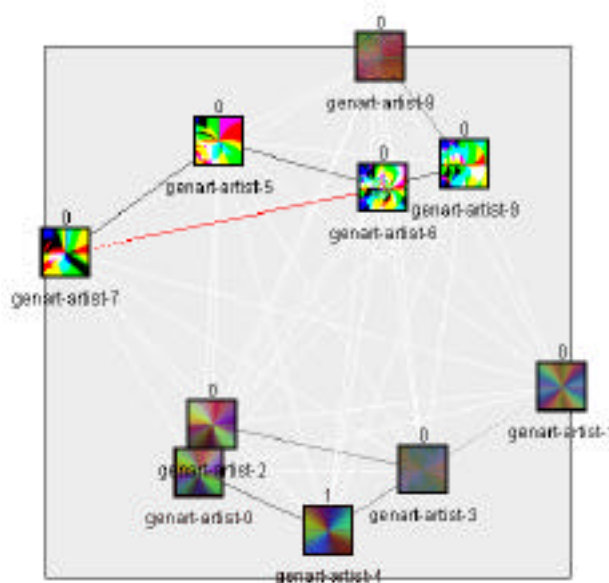


Figure 10: A screenshot of the simulation clearly showing the two cliques. The squares represent agents. The images show the currently selected genetic artwork for each agent. The number above each square shows the agent's attributed creativity. The dark lines between agents indicate the communication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



This experiment has shown that when a group of individual contains mutually incompatible subgroups, the agents in those subgroups will form cliques that communicate credit frequently amongst themselves but rarely to outsiders.

The stability of these cliques depends upon how similar the individuals in different cliques are to one another and how often the agents in one clique are exposed to the artworks of another clique. Other factors that affect judgements of interestingness can have a similar effect, including the type of image processing done in the sensing of images and the processes used to determine the novelty of an image.

## 6 Discussion

In the above simulations the consensus of what is creative has been demonstrated to be a function of both the individual's drive for novelty and the collective knowledge of the group of agents. The definition of a creative artwork is thus a social construct of more than one agent and creativity is an honorary term given to agents that consistently produced artworks appreciated by other agents. Although greatly simplified, this model follows in the footsteps of the systems view of creativity proposed by Csikszentmihalyi (1988, 1999) and is similar to the dual generate-and-test model of social-cultural creativity proposed by Liu (2000).

By studying the emergence of social structures during creative development we hope to gain a better understanding of creativity and the affects that the search for novelty has upon the processes of art and design. Figure 11 displays a small gallery of images recorded as examples of interesting artworks by groups of agents with different hedonic functions.

### 6.1 Future Work

Interactive evolutionary systems like those used by the agents in this research have the potential to create an endless supply of artworks that can quickly overwhelm a user trying to assess all of the possibilities. Future creative support tools might use groups of curious agents like those described above to reduce the information overload of users by only presenting images that are collectively interesting artworks to the artist.

In this respect our goals of future systems are similar to those of Baluja et al. (1994). Baluja, Pomerleau and Jochem attempted to create an evolutionary art system that could ~~to~~ learn the aesthetic preferences of a user by observing the artworks that they selected. However, their

system had limited success and failed to reliably produce new artworks of interest to the user.




Figure 11. A small gallery of artworks produced by agents with different preferences for novelty (N) ranging from N=0 to N=19. In each case a 3-agent group were given a prototype artwork to seed their evolutionary searches. The prototype used was the same for all of the agents in every case.

In contrast, we recognise that it is the search for novelty that is constant in creative activity, not the user's preferences. One possibility is to develop on-line communities comprising of both people and agents interacting through the mechanisms described above. Creative agents would be rewarded for producing interesting artworks while (computational) agents that **failed** to produce interesting artworks would be mercilessly killed off to make way for new agents. In this way, users of the system could create loyal and useful cliques of collaborators. The user would become an elevated peer in a culture of curious agents, obeying and enforcing Martindale's Law of Novelty.

## 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 Hsien-Hui (Michael) Tang for his valuable insights and helpful comments.

## References

- S. Baluja, D. A. Pomerleau, and T. Jochem. Towards automated artificial evolution for computer generated images. *Connection Science*, 6(2-3):325–354, 1994.
- D. E. Berlyne. *Aesthetics and Psychobiology*, Appleton-Century-Crofts, New York, 1971.
- S. Blackmore. *The Meme Machine*, Oxford University Press, Oxford, 1999.
- M. A. Boden. *The Creative Mind: Myths and Mechanisms*, Cardinal, London, 1990.
- S. Colton, A. Bundy, and T. Walsh. On the notion of interestingness in automated mathematical discovery, *International Journal of Human Computer Studies*, 53(3), 2000.
- M. Csikszentmihalyi. 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, 1988.
- M. Csikszentmihalyi. 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, 1999.
- R. Dawkins. *The Selfish Gene*, Oxford University Press, Oxford, 1976.
- R. Dawkins. *The Blind Watchmaker*. Norton, New York, 1986.
- R. Dawkins. The evolution of evolvability. In C. G. Langton (ed.) *Artificial Life*, Addison-Wesley, MA, 1989.
- L. Gabora. The origin and evolution of culture and creativity. *Journal of Memetics: Evolutionary Models of Information Transmission*, 1(1): 1–28, 1997.
- L. Gabora. The beer can theory of creativity. In P. Bentley & D. Corne, (eds.) *Creative Evolutionary Systems*, Morgan Kauffman, 2000.
- T. Kohonen. *Self-Organization and Associative Memory*, 3rd ed., Springer, Berlin, 1993.
- D. B. Lenat. *AM: An artificial intelligence approach to discovery in mathematics*, Ph.D. thesis, Stanford University, 1976.
- P. Langley, H. A. Simon, G. L. Bradshaw, and J. M. Zytkow. *Scientific Discovery: Computational Explorations of the Creative Processes*, MIT Press, Cambridge, MA, 1987.
- Y. T. Liu, Creativity or Novelty? *Design Studies*, 21:261–276, 2000.
- S. Marsland, U. Nehmzow, and J. Shapiro, A real-time novelty detector for a mobile robot. In *Proceedings of EUREL European Advanced Robotics Systems Conference*, Salford, England, 2000.
- C. Martindale. *The Clockwork Muse*, Basic Books, New York, 1990.
- R. Saunders and J. S. Gero. Designing for interest and novelty: Motivating design agents. *CAAD Futures 2001*, to appear.
- J. Schmidhuber. 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, 1991.
- A. Silberschatz and A. Tuzhilin. What makes patterns interesting in knowledge discovery systems, *IEEE Transactions on Knowledge and Data Engineering* 8(6):970–974, 1996.
- K. Sims. Artificial evolution for computer graphics, *Computer Graphics*, 25(4):319–328, 1991.
- S. Todd, and W. Latham. *Evolutionary Art and Computers*, Academic Press, London, 1992.
- M. Witbrock, and S. N. Reilly. Evolving genetic art. In Bentley, P. J. (ed.) *Evolutionary Design by Computers*, Morgan Kaufman, San Francisco, CA, pp. 1–73, 1999.

This is a copy of the paper: Saunders, R and Gero, JS (2001) The digital clockwork muse: A computational model of aesthetic evolution, *AISB'01*, University of York, York, UK, pp. 12-21.