

How to Study Artificial Creativity

Rob Saunders
rob@robsaunders.net

John S. Gero
Key Centre of Design Computing and Cognition
University of Sydney
Sydney, NSW 2006, Australia
+61 2 9351 2328
john@arch.usyd.edu.au

ABSTRACT

In this paper, we describe a novel approach to developing computational models of creativity that supports the multiple approaches to the study of artificial creative systems. The artificial creativity approach to the development of computational models of creative systems is described with reference to Csikszentmihalyi's systems view of creativity. Some interesting results from studies using an early implementation of an artificially creative system, The Digital Clockwork Muse, are presented. The different studies show how the artificial creativity approach supports the study of creativity from a variety of standpoints that mirror the disciplines that study human creativity. The use of artificial creativity allows these different studies to be conveniently conducted on the same computational model and integrated in to a more complete picture of the creative process.

Categories & Subject Descriptors: I.6.3 [Simulation and Modeling]: *Applications*---design computing; I.6.5. [Simulation and Modeling]: *Model Development*---multi-agent systems; J.5 [Arts and Humanities]: *Fine Arts*---evolutionary art.

General Terms: Experimentation.

Keywords: Creativity studies, curiosity, multi-agent simulation.

INTRODUCTION

Artificial creativity [12] is a computational approach to studying creative behaviour using closed-world simulations of social creative systems. In a similar way to Artificial Life, the aim of artificial creativity is to provide insights into the nature of *creativity-as-it-is* by studying *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 can be computationally modelled in artificial societies of agents.

The artificial creativity approach allows researchers to study aspects of creative behaviour at different levels of the creative system within a single model. Artificial creativity systems provide closed-world simulations of creative systems that allow an experimenter to control the details of the simulation that affect

production, communication and recording of ideas and artefacts. This ability permits the study of many configurations not possible in the real world.

As with artificial life, some of the most interesting possibilities supported by artificial creativity systems come from the ease with which researchers can re-run a simulation with different starting conditions or parameters to study how the course of events are changed. These types of experiments are impossible to do in the real world and can potentially tell us a great deal about the nature of creativity.

Artificial creativity is compatible with several different approaches to the study of human creativity, e.g. cognitive science, sociology and history. Mapping these different approaches to the computational realm so that they can be applied to the study of artificially creative systems provides a number of important ways to compare and contrast artificial and human creative systems. The example studies given in this paper give an indication of some of the studies that may be possible in future models of artificial creativity.

The following section provides some background regarding the study of creativity in human creative systems. The artificial creativity approach is then described and some results are presented of studies of artificial creativity using different approaches related to the study of creativity in human creative systems.

CREATIVITY

The apparent need to define the nature of creativity has haunted most attempts to develop models and theories of the processes involved. The difficulty of this task is clear from the number of definitions that can be found in the literature – Taylor [14] gives some 50 definitions. Some researchers have concluded that trying to develop a single definition of creativity is a fruitless task and have looked for ways to conduct their research without the need for a formal definition.

Gardner [4] accepts the judgement of historical record as sufficient grounds to determine the creativity of individuals across multiple domains. He argues that this is the only reasonable way to judge creativity while admitting that it is flawed because it ignores the creativity of individuals that fails to come to the attention of a wide audience. As such, Gardner's approach to determining the creativity of an individual is biased towards those that gain a certain level of fame for their achievements.

Boden [2] provides two definitions of creativity, *psychological creativity* (P-creativity) and *historical creativity* (H-creativity). Within this framework Boden admits that the term creativity has two important meanings in the research community; the first is a label used to identify the processes employed by a creative

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

C&C'02, October 14-16, 2002, Loughborough, Leic, United Kingdom.
Copyright 2002 ACM 1-58113-465-7/02/0010...\$5.00.

individual, the second is a title awarded by society to honour the importance of creative works or creative individuals.

In “The Clockwork Muse” [9] Martindale presented an extensive investigation into the role that individual novelty-seeking behaviour played 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.

Csikszentmihalyi’s Systems View of Creativity

Csikszentmihalyi [3] developed his systems view of creativity after turning his attention away from the question “What is creativity?” and instead asking, “Where is creativity?” The systems view of creativity was developed as a model of the dynamic behaviour of creative systems. Creative systems extend beyond any particular creative individual and include the socio-cultural context within which the individual works. Csikszentmihalyi identified three important components of a creative system; firstly there is the *individual*, secondly there is a social, or interactive, component called the *field*, and thirdly there is a cultural, or symbolic, component called the *domain*. A map of the interactions between these components is given in Figure 1.

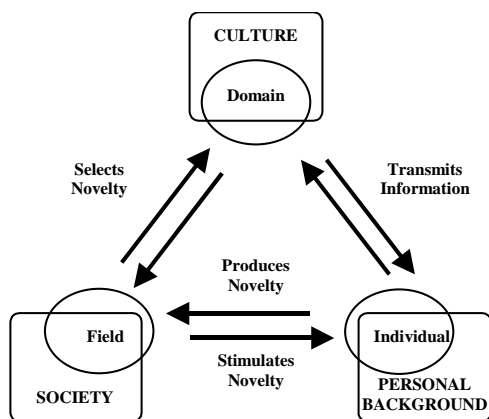


Figure 1. Csikszentmihalyi's systems view of creativity.

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.

Approaches to Studying Creativity

Traditional approaches to studying creativity have tended to focus upon one of these components and acknowledge the other components as context to their subject. This is a natural consequence of the specialised nature of the studies involved, only experts in the disciplines involved can conduct the experiments and few are experts in more than one approach to the study of creativity. For example, cognitive studies have naturally focussed upon the individual and paid less attention, if any, to the social and cultural contexts that influence the creative individual at work. Similarly, sociological studies of scientific or artistic fields provide insights into the nature of the field but rarely attempt to link these back to the cognitive processes of the individuals. Historical studies of particular creative periods rely on the records and artefacts kept by a domain and have little access to the social and cognitive processes that help produce them, other than the records of social structures and the noted findings of earlier researchers. The components of interest in some of the different approaches to studying of creativity are summarised in Table 1.

Table 1. Different approaches to studying creativity and their associated component in the systems view of creativity.

	<i>Individual</i>	<i>Field</i>	<i>Domain</i>
<i>Approaches to studying creativity</i>	Cognitive science	Sociology	History
	Psychology	Anthropology	Historiometric
	Psychometric		Philosophy

ARTIFICIAL CREATIVITY

Artificial creativity attempts to provide a methodology for studying all three components of creative systems using computational models in much the same way that artificial intelligence attempts to provide insights about cognitive processes and artificial life attempts to further our knowledge of biological systems. As such it is concerned with the development of models of creative systems as defined by Csikszentmihalyi rather than the development of cognitive models as it more often the case in artificial intelligence. The essential requirements of a computational model of artificial creativity are:

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

Many of the requirements for artificial creativity are similar to those for computational models of artificial life [7] with the appropriate substitution of a physical environment with a cultural one. The exception is the final requirement of artificial creativity: all agents must be able to evaluate the creativity of artefacts and other agents. No corresponding rule exists in models of artificial life and this is the most important difference between artificial

creativity models and other multi-agent simulations of social processes. The independence of creativity judgements required of agents within an artificial creativity simulation is one of the key aspects of these models that allows them to model interesting features of creative systems. As a consequence of the requirements of artificial creativity social judgements of creativity emerge from the interactions of individuals rather than being dictated by global rules or from outside the simulation by a human experimenter.

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.

The Importance of Emergence

The requirements of artificial creativity are 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 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.

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 are considered emergent properties of creative systems.

THE DIGITAL CLOCKWORK MUSE

The Digital Clockwork Muse is an artificial creativity system developed to explore the role that an individual's search for novelty plays in socially situated creative systems. The Digital Clockwork Muse consists of multiple agents within a single field conducting searches for interesting and potentially creative "genetic artworks" [16]. An example genetic artwork is illustrated in Figure 2, this artwork was evolved by a human user as part of the International Interactive Genetic Art project, for more information on this project see Whitbrock and Reilly [16].

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. Artefacts communicated between agents that are considered interesting by

the peers of its creator are added to the domain as creative artworks. Individuals that produce artworks that are considered creative by other agents are rewarded with "creativity credit".




Figure 2. An example of a genetic artwork.

More detailed and technical accounts of The Digital Clockwork Muse can be found elsewhere [10] the remainder of this document will focus on how the artificial creativity model allows different approaches to the study of creative behaviour during simulation runs.

CURIOSITY, CLIQUES AND COMPLEXITY

The Digital Clockwork Muse provides opportunities to study the effects of the search for novelty on artificial creativity at the level of the individual, the field and the domain. These studies are roughly equivalent to the cognitive, sociological and historical studies that are familiar from studies of human creativity and each provide different insights into the nature of the creative processes implemented.

Curious Individuals

Individual agents in The Digital Clockwork Muse had to be given the ability to judge the potential creativity of artworks as they were produced. To achieve this they needed to be able to independently assess the novelty of an artwork as they encountered it. In addition, the agents needed to be able to use these assessments to search for more novel artworks.

Curious design agents [11] were used to model the novelty-seeking behaviour required for individuals in The Digital Clockwork Muse. A curious design agent embodies a model of curiosity that uses a neural network called a *novelty detector*. A novelty detector can determine the novelty of a new input with respect to all of its previous inputs as a function of the errors generated when it attempts to classify the new input. Using a novelty detector, curious design agents are able to determine the novelty of new artworks as they are produced.

A 32x32-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 fairly complex artworks to be evolved. A 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. A Self-Organising Map (SOM) containing a lattice of 6x6 map neurons is used as the basis for a novelty detector providing a memory 36 image prototypes.

The novelty of each new image is measured as the distance between it and the nearest matching prototype. The distance is

defined as the Euclidean distance between the vectors representing the new image and the closest matching prototype in the 1024 dimensional input space. The novelty values reported in the remainder of this section 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.

The model of curiosity used by the curious design agents in The Digital Clockwork Muse also incorporates a “hedonic function” that transforms novelty into a measure of *interestingness*. The hedonic functions used in The Digital Clockwork Muse are based on the Wundt Curve that Berlyne [1] used as a model for the typical reactions that animals and humans display in the presence of novel situations. The Wundt Curve is illustrated in Figure 3 as the combination of a reward and punishment functions.

Using Wundt Curve hedonic functions the curious design agents in The Digital Clockwork Muse all favour artworks that are *similar-yet-different* to those that have been seen before, however, the agents differ in how similar a new artwork must be for it to be considered highly interesting and therefore potentially creative. The preferred novelty of each agent is expressed as a value N that indicates the amount of novelty associated with peak interest in the Wundt Curve. In The Digital Clockwork Muse, N ranges from 0 to 32; this is equal to the range of the potential classification error generated by the novelty detectors used.




Figure 3. The Wundt Curve: the hedonic function used to calculate interestingness.

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 ones 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 2.

Table 2. 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

Figure 4 shows how the network of communication links that has developed between agents that communicate artworks and evaluations on a regular basis excludes the two agents with different hedonic functions. In the screenshots of the running simulation the squares represent agents; the images in each square shows the currently selected genetic artwork for that agent, the number above each agent shows its attributed creativity, and the lines between agents indicate the number of rewarded communications between pairs of agents.

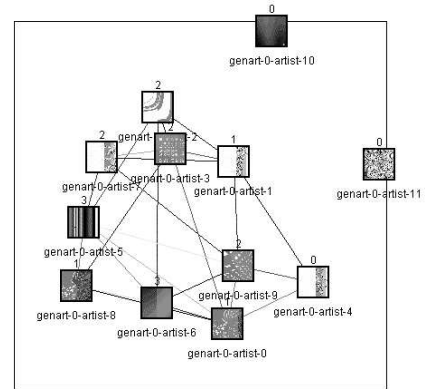


Figure 4. Screenshot of a simulation demonstrating the emergence of the Law of Novelty.

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 a preference for low amounts of novelty, nor agent 11, with a preference for high degrees of novelty, 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.

To better understand the effects of an agent having a different hedonic function to the majority of agents in a population a series of similar simulation runs were performed where the difference between the majority preference for novelty and the two renegade agents is varied from 8, as in the current experiments giving N=3 and N=19, and 1, by giving the two agents hedonic functions favouring N=10 and N=12. The attributed creativity to the agents favouring high and low levels of novelty are shown in Figure 5. The figures plotted against the hedonic are the creativity attributed to an agent relative to the average creativity of the majority of agents that share the same hedonic function.

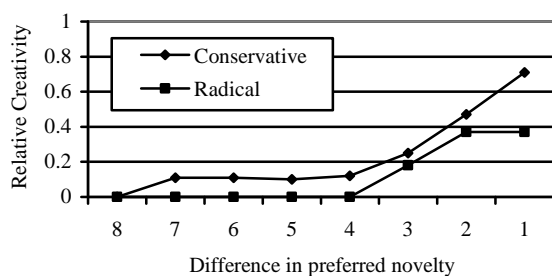


Figure 5. Relative creativity for a range of conservative and radical agents over a range of hedonic values.

Figure 5 shows that attributed creativity varies non-linearly with the difference between an agent’s preference for novelty and the majority. It also shows a slight preference for the works of the conservative agent over the radical one.

The results of this experiment appear to confirm Martindale’s hypothesis generalises to the case where works that are very similar to ones previously experienced are ignored just as much as those that are exact replicas. To avoid being ignored an agent must produce some significant novelty that sets a work apart from previous examples.

The results also indicate that while an agent must produce novelty to be considered creative, it must do so at a pace that matches its audience. There is no advantage in producing many highly novel works if the audience cannot appreciate them. In the first run of the experiment, 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. Indeed, it appears from the series of experiments shown in Figure 5 that erring on the side of caution may be more beneficial than innovating too quickly but more work needs to be done to confirm this experimentally.

The study of the individual in The Digital Clockwork Muse is similar in several ways to the studies of creative humans conducted by cognitive scientists, psychologist, and psychometricians. The design of the individuals follows the traditional approach taken in cognitive science and artificial intelligence of identifying potentially important cognitive functions, in this case novelty detection, interestingness

judgement and curiosity, and then implementing these within a computational model.

The validity of a computational model of cognitive processes is often tested by comparing the behaviour of the model against observations of human subjects. In this case, the validity of the model was tested by comparing the behaviour of a curious design agent against observations about individuals engaged in creative fields, i.e. Martindale’s observations of the importance of the search for novelty.

Continuing the investigation into the relationship between attributed creativity and the deviation of an individual’s preferred novelty from the mode this study provided some quantitative results into the relationship between an agent’s curious “personality” and their creativity, reminiscent of psychometric approach to the study of creative individuals.

Fields 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 6 shows the payments of creativity credit between the agents in recognition of interesting artworks sent by the agents.

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

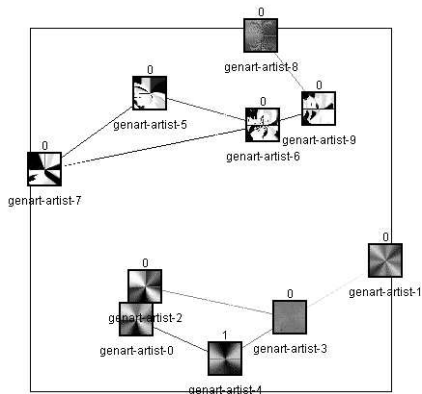


Figure 7. A screenshot of a simulation clearly showing two non-communicating cliques.

Figure 7 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). A brief description of the calculation of fractal dimension used in these studies is given below.

A second pair of groups was simulated with more similar hedonic functions that favoured N=9 and N=12. The communications of credit between agents is illustrated in Figure 8. The results show that while the cliques still form and communication of credit is still concentrated within these cliques, there are more inter-clique communications than before.

An interesting observation about the nature of the communication between cliques can be made from looking at Figure 8 which shows that most of the payments between cliques came from the second group with preference for N=12; only one inter-clique payment was made by a member of the more conservative group that preferred N=9, i.e. between agent-1 and agent-5. This observation is consistent with the earlier observation that it is better to be too conservative than too radical when trying to gain the recognition of others with different preferences for novelty.

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

Figure 8. The communication of credit between two groups of agents having preference for novelty values N=9 and N=12.

There are at least two possible explanations for this observation. The first is that agents with a higher preference for novelty can find the images produced by more conservative agents novel in comparison to the work of their fellow clique members. The second is that agents that prefer lower levels of novelty cannot appreciate the work of more radical agents and hence never attribute any credit to them. It is unclear from these results which explanation is more likely as either would explain the data. Further work may find that both behaviours play a role in the formation of cliques and the unequal communication of credit between them.

The results of this experiment show that when a population of agents contains subgroups with different hedonic functions, the agents in those subgroups form cliques. The agents within a clique communicate credit frequently amongst themselves but rarely to outsiders. 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. Further research is needed to determine whether other factors that can affect judgements of interestingness can similarly affect the social structure.

The studies of clique formation in the fields modelled by The Digital Clockwork Muse provide an indication of how the methods of anthropology and sociology can be applied to artificially creative systems. As a consequence of these studies we can begin to understand how barriers form between different members of a field. The utility of this approach can be seen in the development of the fields of computational sociology and computational anthropology to investigate social phenomena. Potentially, similar models may be able to illuminate issues surrounding the emergence of “paradigm shifts” as documented by Kuhn [6].

Domains of Complexity

To investigate the relationship between the search for novelty and the complexity of the resulting artworks an experiment was conducted to compare agents with different preferences for novelty encoded in their hedonic functions. To measure the complexity of the images the fractal dimension of selected images was calculated. The calculation was performed on the images after image processing to determine the dominant edges so that the fractal dimension would be that of the images as perceived by the agents. The fractal dimension was estimated using the box counting method – this is the same method that Taylor et al. [15]

used to determine the fractal dimension of Jackson Pollock’s drip paintings.

For any two-dimensional image, a measure of its fractal dimension will produce a value between 0.0 and 2.0, depending on how much of the space is filled in the image at different levels of detail. To calculate the fractal dimension of an image a series of grids are placed over the image and the number of boxes occupied by the feature of interest in the image is counted. Figure 9 illustrates the process where the edge segments are the feature of interest.




Figure 9. The box-counting method of estimating fractal dimension of an image.

The fractal dimension can be calculated manually by plotting the count of boxes containing features against the number of boxes per side on a log-log graph and performing a linear regression. The gradient of the line produced is used as an estimate of the fractal dimension. More information about the box-counting method of fractal dimension estimation can be found in Mandelbrot [8].

To investigate the relationship between the preferred degree of novelty and the fractal dimension of the resulting images, two types of agents were used. One type preferred novelty values of $N=18$ and the other type favoured novelty values of $N=11$. Three agents of each type were allowed to explore the space of genetic artworks for 50 time steps.

Figure 10 shows how the average fractal dimension of the images selected by the three agents in each test group changed over time. The graph shows that agents with a preference for greater novelty produce images with higher fractal dimensions, appearing to confirm Martindale’s hypothesis that the search for novelty promotes increased complexity over time [9]. To confirm this relationship between fractal dimension and preferred novelty, similar tests (3 agents/group for 50 time steps) were performed for a total of 19 different test groups with hedonic functions that favoured novelty values in the range $1 \leq N \leq 19$.

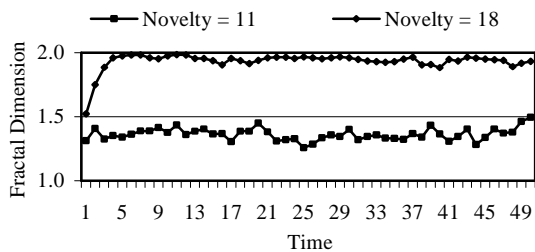


Figure 10. The development of two distinct styles of images with different fractal dimensions in two groups of agents with hedonic functions that peak for the values of novelty indicated.

Figure 11 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. 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.

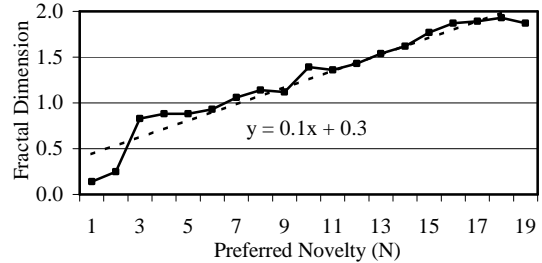


Figure 11. A comparison of the average fractal dimension against a range of peak hedonic values.

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. Figure 12 displays a small gallery of images recorded as examples of interesting artworks by the test groups with preference for the novelty.

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 between complex images than between simple ones.

Analysing the domains produced by fields of curious individuals we have been able to gain some interesting insights into the relationship between the search for novelty and the complexity of the works produced. Analysing the history presented by the domains of small groups of individuals provided clear evidence that the complexity of works increases over time as the individuals explore the space of possible artefacts. It also suggested that agents with different levels of preferred novelty produced artefacts with different levels of complexity. An approach more akin to the multiple domain spanning historiometric approach of Gruber [5], Simonton [13], and Gardner [4] provided a clearer picture of the relationship and a clear link between the preferred novelty of curious design agents and attributed novelty within The Digital Clockwork Muse was developed.




Figure 12. A small gallery of artworks produced by agents with different preferences for novelty (N) ranging from N=0 to N=19.

CONCLUSIONS

The aim of this paper has been to show that the artificial creativity approach to developing computational models of creative systems provides useful basis for a wide range of studies into the three essential components of any creative system. The ability to experimentally study the computational processes of the individual as well as the emergent social structures in the field and the records kept in the domain provides an exciting opportunity to conveniently combine multiple approaches to the study of creativity that are, by necessity, conducted in relative isolation in the real world.

ACKNOWLEDGMENTS

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 much 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.

REFERENCES

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

2. Boden, M. A. *The Creative Mind: Myths and Mechanisms*, Cardinal, London, 1990.
3. Csikszentmihalyi, M. Society, culture, and person: a systems view of creativity, in R. J. Sternberg (ed.), *The Nature of Creativity*, Cambridge University Press, Cambridge, 325–339, 1988.
4. Gardner, H. *Creating Minds*, Basic Books, New York, 1993.
5. Gruber, H. E., and Davis, S. N. Inching Our Way up Mount Olympus: The Evolving System Approach to Creative Thinking, in R. J. Sternberg (ed.) *The Nature of Creativity*, Cambridge University Press, New York, 243–270, 1988.
6. Kuhn, T. *The Structure of Scientific Revolutions*, University of Chicago Press, Chicago, 1970.
7. Langton, C. G. Artificial life, in C. G. Langton (ed.) *Artificial Life*, Addison–Wesley, Reading, 1–47, 1989.
8. Mandelbrot, B. B. *The Fractal Geometry of Nature*, W. H. Freeman & Co., New York, 1977.
9. Martindale, C. *The Clockwork Muse*, Basic Books, New York, 1990.
10. Saunders, R and Gero, J. S. The digital clockwork muse: A computational model of aesthetic evolution, in G. Wiggins (ed.) *Proceedings of the AISB'01 Symposium on Artificial Intelligence and Creativity in Arts and Sciences*, University of York, York, UK, 12–21, 2001.
11. Saunders, R. and Gero, J. S. A curious design agent, in J. S. Gero, S. Chase and M. Rosenman (eds), *Proceedings of CAADRIA'01*, Key Centre of Design Computing and Cognition, University of Sydney, 345–350, 2001.
12. Saunders, R. and Gero, J. S. Artificial Creativity: A Synthetic Approach to the Study of Creative Behaviour, in *Proceedings of Computational and Cognitive Models of Creative Design V* (Heron Island, Queensland, December 2001), Key Centre of Design Computing and Cognition, University of Sydney, 113–139.
13. Simonton, D. K. *Scientific Genius*, Cambridge University Press, New York, 1989.
14. Taylor, C. W. Various approaches to the definition of creativity, in R. J. Sternberg (ed.) *The Nature of Creativity*, Cambridge University Press, 99–124, 1988.
15. Taylor, R. P., Micolich, A. P. and Jonas, D. Fractal analysis of Pollock's drip paintings, *Nature*, 399:422, 1999.
16. Witbrock, M. and Reilly, S. N. Evolving genetic art, in Bentley, P. J. (ed.) *Evolutionary Design by Computers*, Morgan Kaufman, San Francisco, CA, 1999.