A Model of Runaway Evolution of Creative Domains

Oliver Bown
Design Lab, University of Sydney,
NSW, 2006, Australia
oliver.bown@sydney.edu.au

Abstract

Creative domains such as art and music have distinct properties, not only in terms of the structure of the artefacts produced by in terms of their cultural dynamics and relation to adaptive functions. A number of theories have examined the possibility of functionless cultural domains emerging through a runaway evolutionary process. This includes models in which engaging in creative domains is actually counterproductive at the individual level, but is sustained as a behaviour through an evolutionary mechanism. I present a multi-agent model that examines such an evolutionary mechanism, derived from these theories.

Introduction

The study of computational creativity involves both general theory and domain-specific theoretical and experimental studies. Domains such as music, visual art and humour have very different properties owing mainly to the ontological and structural nature of the artefacts produced. But we also know that these domains have different socio-cultural natures. For example, Hargreaves and North (1999) and Huron (2006), discuss social functions and contextual factors that appear to be specific to music, and may not have any relevance to art or humour (although they could). A major contribution to computational creativity therefore involves the computational modelling of specific domains, as in the classic examples described in Miranda, Kirby, and Todd (2003), and more abstract notions of creative domain dynamics, as studied by Saunders and Gero (2001) and Sosa and Gero (2003), drawing on the theoretical formulation of Csikszentmihalyi (1990). The specific analysis of creative domains – their origins, dynamics and relations to individual motivations – makes a critical contribution to computational creativity by framing how we should understand the evaluation of automated creative agents acting in those domains. This paper follows the latter work but looks at the more fundamental evolutionary question of the emergence of creative domains, i.e., how humans came to exhibit behaviour in specific realms such as art and music, either through genetic or cultural evolutionary processes.

The approach used here follows the epistemological method, established in multi-agent modelling fields such as artificial life (Di Paolo, Noble, and Bullock, 2000) and computational social science (Conte et al., 2012), of attempting to reveal novel mechanisms through the study of the emergent qualitative outcomes of local interactions in computer simulations.

The model presented in this paper is based on theories of the evolution of music and takes the form of a minimal abstract model of biological evolution. However, it does not directly look at modelling music, but at a proposed model of underlying social interactions that would allow a runaway evolutionary process to take place. Theoretically this is grounded in the ideas of cultural evolution provided by Boyd and Richerson (1985) and Laland, Odling-Smee, and Feldman (2000). In the language of Laland, Odling-Smee, and Feldman (2000), the model is an experimental study of the 'construction of cultural niches' which remains generic for the sake of simplicity, but could be later developed into a specific model of the construction of a music niche, or applied comparatively to different creative activities. A niche is defined here as a site of fitness acquisition for an individual. Niches can be pre-existing, as in the use of trees for birds, or constructed, as in the alteration of an ecosystem by a beaver building a dam. The model can be interpreted as a general model of runaway evolution of creative domains.

In my conclusion, I discuss the applicability of this model, and more generally this type of modelling, to developing a richer understanding of creative domains that may inform computational creativity. This and similar models provide candidate properties of creative domains that directly inform the way we view the analysis and evaluation of individual creative systems within specific domain contexts.

'Runaway' Theories of the Origins of Musical Behaviour

The origins of music are mysterious and highly contested. In *The Descent of Man* (1883), Darwin introduced the principle of sexual selection and suggested that various aspects of human appearance and behaviour, including music, may be sexually selected. The theory of sexual selection states, in modern genetic terms, that since reproductive achievement is key to the perseverance of genetic lineages, then genetic adaptations that increase ones attractiveness to potential mates will prosper. The theory of sexual selection was developed considerably by Fisher (1915), who proposed that a runaway selection of arbitrary traits could occur if male traits and female preferences coevolved (since females typically have the greater investment in reproduction they are typically the choosier sex). The question of whether sexually selected traits can be fully arbitrary has been the subject of much debate. As part of a general principle that underlies the contemporary study of 'honest signalling theory', Zahavi (1975) proposed that female preference is likely to be guided towards traits that are actually an external (visible or audible) indicator of some positive quality. Thus when male traits and female preferences coevolve, it is those pairings that lead to stronger fitter males that persevere. For example, the quality of a bowerbird's nest indicates the ability of the bowerbird in foraging.

More recently, Miller (2000) has revived the argument that music, amongst other aspects of human appearance and behaviour, is sexually selected. Miller presents musical ability as an indicator trait of general intelligence and health. The theory continues to attract attention but competes with a number of other theories about the origins of musical behaviour. Two strong competing theories are that music serves some cooperative function (Brown, 2007), and that music has no function at all, instead being a cultural innovation that exploits human aesthetic preferences (Pinker, 1998). Both runaway sexual selection and this cultural exploitation theory fit well with an apparent lack of function in music. Although evidence does exist to support social functions in music that would support the cooperative view, this view has also struggled to gain traction due to uncertainty surrounding plausible mechanisms for the evolution of altruism (Fisher, 1958). The sexual selection view has also been criticised because of a lack of typically sexually dimorphic traits in humans with respect to music, and the prevalence of music in situations that appear to have nothing to do with courting, such as at funerals and heavy-metal concerts (Huron, 2001).

However, runaway evolutionary processes are not limited to sexual selection. Zahavi's (1975) examples of honest signalling, for example, extend to other coevolutionary situations. Boyd and Richerson (1985) propose a runaway cultural evolutionary process based on a set of heuristics describing how individuals adopt cultural traits, based on frequency and status. They hypothesise that people are more likely to adopt a cultural trait the more other people adopt that trait, and the higher the status of the people are. They also propose that minimal discrimination is applied to the choice of traits to adopt, on the basis that false positive assumptions are more acceptable than false negative assumptions. In this way potentially arbitrary traits exhibited by high status individuals can easily and rapidly become adopted. Blackmore (1999) develops similar principles through the theory of memetics, and suggests that various aspects of culture, even language, might be understood as having emerged as 'parasites', exploiting human behaviour to become established. These views align with Pinker's view of music as a functionless cultural innovation. Such theories also raise the possibility of a coevolution between genes and culture, which has been explored by a number of theorists, most notably Laland, Odling-Smee, and Feldman (2000). Their extensive theoretical and empirical review suggest that sexual selection and Boyd and Richerson's runaway cultural evolution are just instances of a more general tendency for runaway evolutionary processes to occur between environments and organisms, and that there may be other ways in which runaway evolution could occur in cultural systems. Here the term 'environment' includes culture, and culture is viewed as a site with great potential to exhibit runaway evolutionary processes.

A Model of Runaway Evolution

Little research has been done into how specific cultural forms such as music might be explained by runaway evolution. In this paper, I present a model that provides a very simple mechanism whereby runaway selection of arbitrary cultural domains can become established.

The model is predicated on the broad question underpinning runaway evolutionary processes: under what circumstances will populations of individuals evolve to exhibit traits or engage in behaviour that has no net advantage? Models such as those of runaway sexual selection present such circumstances and show how they are viable. Whilst peacock tails are a burden to peacocks as far as flying or escaping predators are concerned, they give the individual peacock with the better tail a reproductive advantage and thus a net fitness gain. The peacock's tail is understood in terms of the niche created by the peahen's evolved sexual preference, and vice versa. By analogy, in the present case, the goal is to examine examples of cultural behaviours where a similar emergent cultural niche could be established. In our case, we choose to examine a scenario that is not underpinned by sexual selection, but by economics. Primate social organisation is sufficiently complex to lend to the idea that human evolution has been guided by very simple but significant forms of economic interaction. In particular, simple forms of transferrable wealth might have had the capacity to influence fitness dynamics, stimulating the emergence of new cultural niches through positive feedback. Transferrable and cumulative wealth has the capacity to influence evolutionary fitness by allowing one person to effectively take fitness from another person, and, on a macroscopic scale, for societies to develop systems by which to organise their collective wealth, in effect providing some top-down determination of fitness. Under such circumstances, the nature of that social system would have a significant influence on an individual's choice of fitness strategies and this might ultimately have an influence on culturally evolved behaviour, and possibly even a genetic influence. Note that transferrable wealth could mean something such as rights to land that is not achieved technologically, but merely requires a simple concept of ownership or title, although in the present case wealth is also considered cumulative, which might entail something being harvested, or simple things such as clothing being made. Given their simplicity, these factors plausibly predate the creative domains under consideration.

But what has this got to do with creative domains such as art and music? A number of recent studies have looked at how creative success is organised at a social level, suggesting that there is inherent positive feedback in the way that we allocate reward for creative achievements. Salganik, Dodds, and Watts (2006), for example, show that music ratings are directly influenced by one's perception of how others rated the music, not just in the long term but at the moment of making the evaluation. The result is a winner-takes-all outcome, where a piece of music that is rated highly by others is more likely to be highly rated in the future, as long as people are aware of the already-high regard given to the work. Rather than directly appraising creative works in terms of their content, they are appraised as social artefacts, subject to social processes that transcend the creative content itself. If this is true, then one potential effect of individuals engaging in creative domains is to create winner-takes-all redistributions of some social entity, most broadly described as *prestige*, that may be assumed to relate in some way to wealth.

Accepting the assumption that any given creative domain has no other fitness-enhancing function, then in evolutionary terms it can be understood as a time and effort commitment that needs to be explained. The present model looks to reduce such a scenario to its simplest abstraction and consider the evolutionary effects (whether generic or cultural). In particular, it asks whether it is possible that the creative domain acts to reinforce itself over time, thus providing a evolutionary explanation in the form of niche construction. For this to be demonstrated, a population must be shown to transition from not engaging in the creative behaviour to engaging in it. This occurs when those who engage in the creative behaviour are more successful than those who do not engage over evolutionary time. The model presented here looks at how this can happen over evolutionary time, despite the net average benefit for engaging in creative domains being lower than for avoiding them.

Model Design [1]

The model has a very specific purpose, which is to show how an arbitrary activity can emerge amongst a population of rational selfish agents. Underlying the model, a simple economic system is implemented in which wealth is tied to evolutionary fitness. Agents with higher wealth have a greater chance of survival and are therefore driven by natural selection to maximise wealth. The purpose of the model is to demonstrate evolutionary scenarios in which emergent social conditions favour acting in an apparently irrational way, by engaging in an arbitrary functionless behaviour: a 'game'. The functionless behaviour in turn provides the conditions for runaway evolution.

Note that evolution here can refer to the evolution of genes or of culturally (vertically) acquired traits, interchangeably. Thus the model works as either a biological or a cultural evolutionary model. For the purpose of this paper I refer to genes in the model, but these can be replaced by 'memes' that are vertically transmitted.

The model consists of a fixed population of N agents. Evolutionary competition is implemented through tournament selection. Each agent has the following genetic variables:

- Tendency to play the game (G_i): the probability that an agent will chose to play the game in a given round. At each time step, each agent is identified either as a gamer or a non-gamer;
- Competence (C_i): the game is predominantly random, but there is a bias towards agents with a higher competence;
- Taxation Vote (T_i): all agents vote on a level of taxation that non-gamers should pay into the game, the tax at each round is the average of these T_i.

Each agent also has a wealth variable, W, which is modified through transactions as described in the sequence below.

The following sequence is run at each time step:

1. All agents accumulate a fixed 'pay', p.
2. A globally imposed non-gamer 'tax', t is calculated as the average of all agents' Taxation Votes, T_i.
3. All agents are asked if they wish to play the game in the current round, resulting in a number n of gamers. The tendency to play the game, G_i, is treated as a probability that determines this choice.
4. All gaming agents pay a fixed cost, c, whilst all non-gaming agents pay the non-gamer tax, t. Non-gamer agents also receive the fixed non-gamer bonus, b.
5. The game winner is determined as follows: two different agents are randomly chosen from the list of gamers. The agent with the greatest competence, C_i, out of these two candidates wins. In the case of equal ability to cheat, a random agent is the winner. The winner receives all of the bids, $n \times c$, and all of the tax, $(N - n) \times t$.
6. A fixed number m of reproductive tournaments are run as follows: two different agents are randomly chosen from the population. The agent with the greatest wealth is the winner. In the case of equal wealth, a random agent is the winner. The loser is replaced by a child (mutated copy) of the winner. The parent gives a fixed proportion, w, of its wealth to its child.
7. All agents' wealth is depreciated by a wealth depreciation coefficient, d ($0 \leq d \leq 1$). Each agent's wealth is scaled by this number.

Children's G_i, C_i and T_i values, the genetic variables, are copies of the parent with a Gaussian mutation with a standard deviation of 0.001. G_i and T_i values are constrained between 0 and 1. C_i values are only constrained with a lower bound of 0. The parent gives a fixed proportion, w, of its wealth to its child. Unless otherwise stated, initial values for all agents are $G_i = 0$, $C_i = 0.5$ and $T_i = 0$.

The model variables used in the studies detailed in this paper used the values specified in Table 1.

Starting from an initial value of zero, an increase in the mean tendency to play the game, $\bar{G}$ is then interpreted as a scenario in which game-playing behaviour has become established in the population. The model is designed to reveal

---

[1] All code for the software model can be found at https://www.dropbox.com/s/48oy1v32lx0utp0/LotteryMain.java. The variables described in this paper differ from those in the code, which are based on the scenario of a lottery game.

Var	Description	Value
p	Pay for all agents at each time step	1
c	Cost of bid paid by gamers	1
b	Bonus paid to non-gamers	1
m	Reproduction tournaments per iteration	10
w	Proportion of wealth paid to children	0.2
d	Depreciation of wealth at each time step	0-0.999

Table 1: Values used in experiments. All values are fixed except the experimental variable d.

the conditions that are required for this to arise. $\bar{G}$ is subject to dynamic selection pressures and can also drift, if no strong selection is observed. Through propagation through a population the range of drifting G values can appear to have low variance, so low variance is not considered sufficient to indicate strong selection. Constraint of the variable to a specific range over a long period of time and multiple runs is used: if $\bar{G}$ sits consistently above 0.8, it is concluded that a game behaviour has emerged in the population.

We assume that individuals are equally able to generate transferrable wealth at a fixed rate, p, per time step. For the game, players put a unit of their wealth, c, into a pot and one individual, chosen at random, wins the entire pot. In addition, we assume that game-playing has a fixed time cost. This is implemented as a further payment, b, to non-gamers. The relative values of p, c, and b therefore define a space of possible model parameters with possible outputs with regard to how G evolves.

Results

The wealth depreciation coefficient (d) was compared across 4 values, 0, 0.9, 0.99 and 0.999. In the first case, wealth is transitory, acquired at the beginning of each time-step, then either spent in the game or kept, and then used to compete in tournaments. For values of d approaching 1, wealth becomes increasingly cumulative. This has two implications: firstly, wealth reaches higher levels, since with a constant income the stable state wealth value is greater. Greater wealth takes longer to accumulate and means that individual gains are ultimately less relevant. Secondly, the gains of short-term successes stick around longer and are more likely to transform into reproductive success. These can also be transferred to children.

Figure 1 shows model outcomes for the values of d, 0.9, 0.99 and 0.999. Each graph shows the average of the 'tendency to play the game' genetic variable, $\bar{G}$, in the population over time, with 20 runs of the model superimposed on each graph. $\bar{G}$ tends toward its upper bound in models with $d = 0.999$, whereas it does not drift far away from zero in models with low d ($d = 0$ and $d = 0.9$). Even for $d = 0.999$ there is the potential for $\bar{G}$ to drift down as well as up, indicating that population-wide game behaviour under favourable circumstances is not as strong an evolutionary stable-state as game-avoidance under unfavourable circumstances. These results show that the durable, transferrable forms of wealth discovered by humans create a situation conducive to the formation of game-playing.

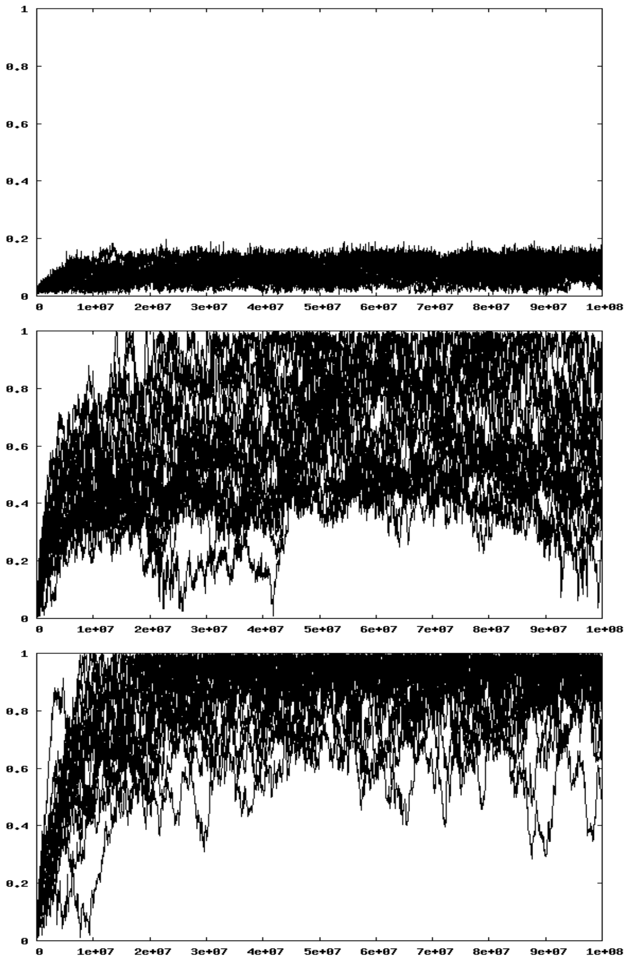

Figure 1: Evolutionary runs with wealth depreciation coefficient, d, values of (from top to bottom) 0.9, 0.99 and 0.999. Each graph shows the mean 'tendency to play the game' genetic variable, $\bar{G}$, evolving over 100 million time-steps, repeated over 20 runs of the model. The taxation vote is allowed to evolve genetically.

Figure 2 shows a typical instance of the model for $d = 0.999$ and evolvable taxation vote T, with $\bar{G}$ in red and the mean Taxation Vote genetic variable, $\bar{T}$ in green. Both values are attracted towards their upper bound of 1, with $\bar{T}$ more inclined to drift. It may be a reasonable assumption that these variables are positively mutually reinforcing, though this has yet to be tested.

In order to understand the specific economic pressures on individuals, a simplified study was conducted with the taxation vote set to a fixed value. To further clarify the model, the accumulated tax was not passed to the game winner, as described above, but was instead discarded. This makes it easier to measure the average expected incomes of individuals in the non-gaming and gaming categories, since average incomes are no longer frequency-dependent (as compared to the standard model where tax channels wealth from non-gamers to the winning gamer). In this simplified model,

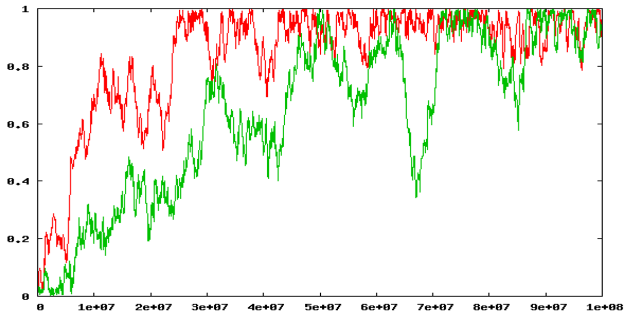

Figure 2: An example evolutionary run with wealth depreciation coefficient, d, of 0.999. The graph shows the mean 'tendency to play the game' genetic variable, $\bar{G}$, in RED, and the mean taxation vote genetic variable, $\bar{T}$, in GREEN, evolving over 100 million time-steps for a typical run.

non-gamers gain $(p + b - T)$ units of wealth at each time step. Gamers do not gain benefits b or pay tax T. Since the game is zero-sum their average income is simply p. Non-gamers all receive the average income, whereas gamers' real incomes are skewed according to the outcome of individual games.

Figure 3 shows the emergence of game playing (situations where $\bar{G}$ tends towards 1) for different values of T, under these conditions. For $T = 0.4$ game playing begins to emerge. The transition from non-game to game takes the form of a sudden phase shift with an erratic onset, and no transitions occur in the opposite direction, implying that game-playing is evolutionarily stable in the population once established. With $T = 0.6$ game playing consistently emerges. In the latter case, the average non-gamer income is $(p + b - T) = (2 + 1 - 0.6) = 2.4$ whereas the average gamer income is $p = 2$. Therefore even when the non-gamer group is fitter on average, the gamer group comes to dominate. This shows a minimum requirement for game playing to emerge. By comparison, the graph at the bottom of Figure 1 shows that this result is robust if T is allowed to vary genetically, even when initial values for G and T are zero.

Figure 4 shows the mean competence genetic variable, $\bar{C}$ increasing steadily without limit for the same run as the graph in Figure 2. $\bar{C}$ exhibited this increase consistently across all runs, even with $d = 0$. By the model design there can be no circumstances under which lower C is advantageous, and always the occasional accidental game that selects in favour of higher C. The purpose of modelling C is not to show that it increases, which is inevitable and obvious, but to show that it has no impact on the emergence of the game behaviour, despite undermining its 'fairness'. We can say that random success is sufficient for the game to emerge, and may enable the initial adoption of the behaviour, but that it is not strictly required. What matters is that the game is robust once established, and creates a stable scenario in which C is driven to evolve. In this model, C is just a numerical variable that is driven to evolve indefinitely towards higher values, but in its place more complex models could explore the potential for the game-playing niche to drive a runaway arms-race of game-playing skill, with each winner passing on the greatest skill traits to the next generations.

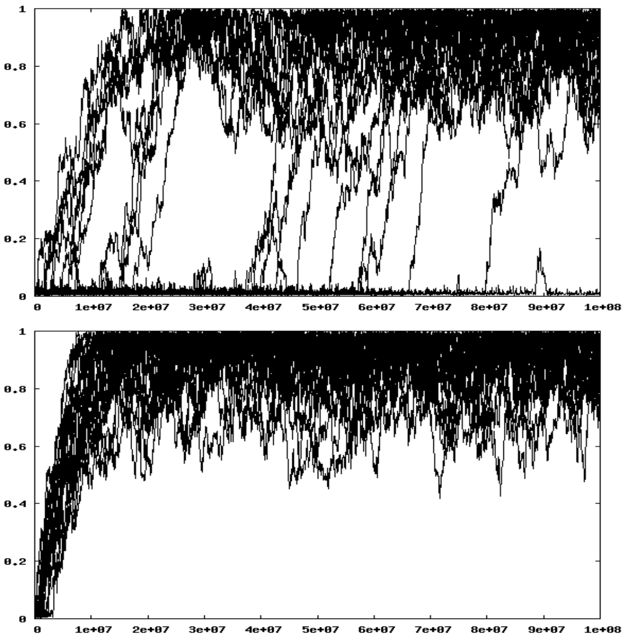

Figure 3: Evolutionary runs with wealth depreciation coefficient, d, values of (from top to bottom) fixed 0.999. Each graph shows the mean 'tendency to play the game' genetic variable, $\bar{G}$, evolving over 100 million time-steps, repeated over 20 runs of the model. In this case the taxation vote is fixed at 0.4 (top) and 0.6 (bottom). Furthermore, in these instances taxes are *not* passed onto the game winner but are simply discarded.

Discussion

Summary of Results

To summarise the key results, the model shows how a population can evolve an apparently economically irrational behaviour that drives inequality. A greater durability of wealth increases the tendency for game playing to occur, even if the net benefit to the average individual is lower. The emergence of evolutionarily stable game playing behaviour creates a selective pressure driving the constant and rapid increase in game playing ability, but as the population evolves together towards greater competence, the game itself is sustained. As discussed, the properties of this system resemble a set of hypothesised properties of creative domains, satisfying a niche construction view of their emergence.

The results therefore reveal a hypothesised emergent cultural niche which, too all extents and purposes, is functionless, but provides a site for individual fitness acquisition (albeit achieved by lottery) by individuals, and drives a runaway competitive coevolution amongst the population of greater competency in this domain.

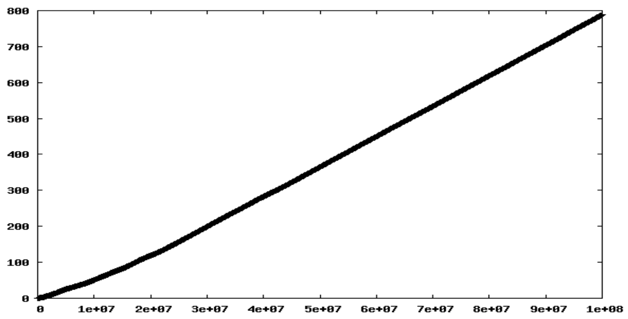

Figure 4: A simulation run ($d = 0.999$) showing the mean competence genetic variable, $\bar{C}$, evolving over time. In all cases, including $d = 0$, $\bar{C}$ increased without limit.

On Randomness

The choice to base the model on a lottery-like game was not discussed in the theoretical background, but is also grounded in a well-founded evolutionary concept. Given the evidence for winner-takes-all processes in human artistic domains, the possibility that randomness is a significant part of the process is actually something that should be seriously considered. A possible role of randomness in structuring social systems, proposed by Wilson (1994), supports a functional role for randomness.

Along with heredity and meritocracy, Wilson (1994) shows that chance can and does play a role in the construction of socially structured systems. The clearest and most striking example of this is the determination of gender, a stochastic process occurring in development, that leads to a prominent social distinction, underpinned by physiological divergence. Looking at the abstract properties of our biological system of gender, Wilson (1994) argues that there may be any number of other behavioural traits determined through a similar process: genetically determined phenotypic variations derived from a common genotype, allocated stochastically. They are, by this definition, not environmentally determined, and are therefore strictly chance allocations, not local adaptations. It is through a stochastic process that a given distribution of possible behaviours emerges, just as in the case of gender, where we end up with a roughly 50-50 split.

Wilson proposes variation along a boldness-shyness personality scale as a candidate example. Assume that boldness and shyness are both proven to be optimal behavioural strategies in different social contexts (in the context of art we could map these onto traits such as creativity and conformity). Typically we think of phenotypic plasticity as the only approach to arriving at good context-dependent behavioural strategies such as these. A plasticity-based view of these traits is that an individual would learn from cues in their environment to be either shy or bold. An equally plausible explanation, Wilson argues, is that the trait is randomly assigned by a stochastic developmental process. Assuming that, to some extent, individuals can find roles that suit their phenotypes (*i.e.*, there are places in the social system where both shy and bold individuals can thrive better than the other), and that an appropriate range of roles is available, then all individuals can emerge well-adapted. Thus a social structure that demands a mix of traits can coevolve with this kind of stochastic allocation of traits. The principle of self-organisation can explain the resulting assignment of roles.

This explanation is also satisfying because the genetic mechanism for stochastically switching between two evolved behavioural variations is arguably simpler than the psychological mechanism required to work out which behaviour strategy is successful in a given, novel context. In addition, the precise source of randomness might be at a number of different stages other than in the genetics. For example, boldness-shyness development could be triggered by events that are effectively random, *i.e.*, there is nothing in the content of the trigger that conveys relevant information about the environment. In the case of creative domains, as suggested by the present model, creative success could be allocated randomly, with the effect that those creatively successful individuals act to reinforce the existence of the creative domain for future generations. This is only to say that random allocation of creative success may be sufficient for the creative domain to work. In reality, creative success may also depend on non-random processes, as with our competence variable.

An important clarification of this principle is that it is not necessary for every individual to do equally well out of the situation for it to be evolutionarily viable: a principle well established by sociobiologists, as in the respective reproductive fitness of different individual ants in a colony. Instead, the process can produce clear inequalities. This parallels the principle of kin selection; kin-directed altruism is able to evolve in proportion to the degree of relatedness between kin, based on the fact that altruism between close kin is as good a way for genes to persevere as individual selfishness. Kin-selection is widely believed to be the most robust mechanism by which cooperative behaviour emerges in nature (Maynard Smith and Szathmáry, 1995).

Conclusion

The model of runaway evolution presented in this paper simply provides a mechanism whereby a pattern of behaviour resembling human creative domains can emerge. The provision of a mechanism does not in any way help prove the theory that music and other creative domains emerged through runaway evolution, but enables predictions derived from the mechanism provided.

The simulation model can be tested against studies of the nature of creative success over multiple generations, taking into account the relationship between creative success, core economic motivations, overall fitness and other contextual factors. In particular, the model predicts that the motivation to engage in creative domains is irrational in the short term but evolutionarily stable in the long term. We can test this by looking at the immediate payoff to art practitioners of varying levels of success. The model predicts that this payoff would be poor in the short term, but that this apparently irrational behaviour could be explained by a process of reinforcement occurring at the social level, whereby creatively successful individuals effectively assert the status quo.

Such factors provide a wider context for thinking about the evaluation of artificial creative systems. Evaluation as presently conducted on an individual case-by-case basis (system by system or output by output) may need to be revised to take into account a more complex understanding of the relationship between long-term creative dynamics and short term creative success. Rather than building one virtual Mozart or virtual Picasso, we may need to deploy millions of them in virtual communities in order to truly understand creative success.

References

Blackmore, S. J. 1999. *The Meme Machine*. New York: OUP.

Boyd, R., and Richerson, P. J. 1985. *Culture and the Evolutionary Process*. Chicago, IL, US: University of Chicago Press.

Brown, S. 2007. Contagious heterophony: A new theory about the origins of music. *MusicæScientiæ* 11(1):3–26.

Conte, R.; Gilbert, N.; Bonelli, G.; Cioffi-Revilla, C.; Deffuant, G.; Kertesz, J.; Loreto, V.; Moat, S.; Nadal, J.-P.; Sanchez, A.; et al. 2012. Manifesto of computational social science. *The European Physical Journal Special Topics* 214(1):325–346.

Csikszentmihalyi, M. 1990. The domain of creativity. In Runco, M., and Albert, R. S., eds., *Theories of Creativity*. Sage Publications. 190—212.

Darwin, C. 1883. *The Descent of Man and Selection in Relation to Sex*. New York, USA: Appleton and Co.

Di Paolo, E.; Noble, J.; and Bullock, S. 2000. Simulation models as opaque thought experiments. In A.Bedau, M.; McCaskill, J. S.; Packard, N. H.; and Rasmussen, S., eds., *Articial Life VII: Proceedings of the Seventh International Conference on Articial Life*, 497–506. Cambridge, MA: MIT Press.

Fisher, R. A. 1915. The evolution of sexual preference. *Eugenics Review* 7:184–192.

Fisher, R. A. 1958. *The Genetical Theory of Natural Section*. London: Dover.

Hargreaves, D. J., and North, A. C. 1999. The functions of music in everyday life: Redefining the social in music psychology. *Psychology of Music* 27(1):71–83.

Huron, D. 2001. Is music an evolutionary adaptation? *Annals of the New York Academy of Sciences* 930(1):43–61.

Huron, D. 2006. *Sweet Anticipation*. MIT Press.

Laland, K. N.; Odling-Smee, J.; and Feldman, M. W. 2000. Niche construction, biological evolution, and cultural change. *Behavioral and brain sciences* 23(01):131–146.

Maynard Smith, J., and Szathmáry, E. 1995. *The Major Transitions in Evolution*. New York: Oxford University Press.

Miller, G. 2000. Evolution of human music through sexual selection. In Wallin, N. L.; Merker, B.; and Brown, S., eds., *The Origins of Music*. Cambridge, MA, USA: MIT Press. 329–360.

Miranda, E. R.; Kirby, S.; and Todd, P. M. 2003. On computational models of the evolution of music: From the origins of musical taste to the emergence of grammars. *Contemporary Music Review* 22(3):91–111.

Pinker, S. 1998. *How the Mind Works*. London, UK: Allen Lane The Penguin Press.

Salganik, M. J.; Dodds, P. S.; and Watts, D. J. 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *science* 311(5762):854–856.

Saunders, R., and Gero, J. S. 2001. Artificial creativity: Emergent notions of creativity in artificial societies of curious agents. In *Proceedings of the Second Iteration Conference*.

Sosa, R., and Gero, J. S. 2003. Design and change: A model of situated creativity. In Bento, C.; Cardoso, A.; and Gero, J., eds., *Approaches to Creativity in Artificial Intelligence and Cognitive Science, IJCAI03*, 25–34.

Wilson, D. 1994. Adaptive genetic variation and human evolutionary psychology. *Ethology and Sociobiology* 15:219–235.

Zahavi, A. 1975. Mate selection—a selection for a handicap. *Journal of theoretical Biology* 53(1):205–214.