

Enhancing Creativity with Niche Construction

Jon McCormack

Centre for Electronic Media Art, Monash University, Clayton 3800, Australia
Jon.McCormack@monash.edu

Abstract

We address the question of how processes from evolutionary biological ecosystems can be abstracted and beneficially applied in creative domains. Evolution is a process capable of generating appropriate (fit) novelty in biological systems, so it is interesting to ask if it can do so in other, non-biological systems. Past approaches have focused on optimisation via fitness evaluation (either machine representable or human evaluated), but this is ill-suited to creative systems, as creativity is not necessarily an optimisation process. Our approach is to consider the creative system as a virtual evolutionary ecosystem, specifically adopting the process of niche construction. We show how the abstracted niche construction process can be applied to an agent-based line drawing system, enhancing the diversity and heterogeneity of drawings produced over a version without niche construction.

Introduction

Two well known systems exhibiting creativity are the human brain and evolution. While advances in neurological understanding of creative processes and aesthetics are ongoing (Perlovsky, 2010; Griffiths, 2008; Ramachandran and Hirstein, 1999), both the cognitive and social processes that lead to creative outcomes remain difficult to quantify, and hence, to simulate. Evolutionary processes, on the other hand, are far better understood and continue to be successfully studied using a variety of simulation methods.

In this paper we explore the adaptation of evolutionary ecological processes to problems in creative design. As a process, evolution is eminently capable of novel design, having innovated things such as prokaryotes, eukaryotes, higher multicellularity and language, through a non-teleological process of replication and selection (Maynard Smith and Szathmáry, 1995; Nowak, 2006). While much exists on what constitutes human creativity – e.g. Boden (2004); Sternberg (1999)) – for the purposes of this paper we consider creativity more generally as the *appropriate novelty* exhibited by a system. ‘Appropriate’ in that the artefacts produced are fit or useful in some domain, and ‘novel’ in

that the system is capable of repeatedly producing artefacts that it has not produced before¹.

Darwinian processes of selection and replication with difference only provide a simplified picture of natural evolution. Many have argued that explaining the growth of complexity that typifies the creativity of evolution requires a broader consideration of the systems of the natural world (Maynard Smith and Szathmáry, 1995; Laland et al., 1999; Gould, 2002). In recent years, that has meant, for example, increasing our understanding of (i) the effects of evolution on the processes of ontogenetic development (Carroll, 2005) (ii) the interdependent relationships between species and their environment: *ecosystems*. This second approach is the one adopted in the work described here.

Evolution and Aesthetic Creativity

The field of Evolutionary Computing (EC) has adopted the metaphor of genetic evolution to successfully solve problems in search, optimisation and learning. Where EC has been less successful, however, is in tackling problems of creativity, in particular artistic creativity, as it is difficult to conceptualise creative artefacts in terms of a single (or multi-objective) optimisation or general machine-representable fitness evaluations.

A popular EC approach to using evolution in creative contexts is the *Interactive Genetic Algorithm* (IGA), in which the fitness evaluation of a standard genetic algorithm is performed by a human, who may use any (subjective) criteria to assign fitness to individuals in a population (Takagi, 2001).

In the context of the application presented in this paper (line drawing) the system of Baker and Seltzer (1994) used variable length genomes representing an ordered set of strokes to define a line drawing. Each stroke included parameters in the genome to affect the way drawing is interpreted, including space enclosing, relation to the next stroke (e.g. separate or joined) and symmetry operations. Drawings were evolved using an IGA. The system could be seeded with random genotypes or genotypes created by interpret-

¹For a more formal specification of this relatively informal definition, see McCormack (2010).




Figure 1: Example organism viability curves for reproduction, growth and survival, from (Begon et al., 2006).

ing the strokes of a human artist. The *Drawbots* system of Bird et al. (2008) attempted to create a line-drawing robot using evolutionary robotics. They defined “implicit” fitness measures that did not restrict the type of marks the robot drawer should make, including an “ecological model” involving interaction between environment resource acquisition and expenditure through drawing. However, the results demonstrated only minimal creativity, and the authors concluded that fitness functions that embodied “artistic knowledge about ‘aesthetically pleasing’ line patterns” would be necessary if the robot were to make drawings worthy of exhibition.

Formalised “aesthetically pleasing” fitness measures of any generality have been difficult to find, despite a number of attempts (see e.g. Birkhoff (1933); Staudek (2002); Ramachandran (2003); Svangård and Nordin (2004); Machado et al. (2008)), hence the use of the IGA. While the IGA has achieved some success in a variety of domains, in general it suffers from a host of problems, particularly for creative applications (McCormack, 2005). The most commonly cited of these is “user fatigue”, where human users quickly tire of the repetitive act of phenotype evaluation (Takagi, 2001), limiting the range of evolutionary exploration possible. In general, IGAs are more valuable to non-experts, who may lack the sophisticated understanding of how to design and manipulate a medium for creative purposes.

More importantly, for most creative domains the idea of evolving towards a single optimum is counterintuitive, as an artist or designer normally produces many new artefacts over their professional lifetime. New designs often ‘evolve’ from previous ones, offspring of both the originating artist and her peers (Basalla, 1998). Indeed, as Basalla (1998) and others have pointed out using the example of technological evolution, the Western emphasis on individual creativity (reinforced socially through patents and other awards) obscures the important roles played in the evolutionary ecosystem of interactions between environment and prior work of many individuals.

Thus, an alternate approach to the narrow individual optimisations of standard EC methods, is to consider the in-

teraction of components in an evolutionary ecosystem, as such a system can potentially exploit facets of evolution other than single optimisations. In the research presented in this paper, we examine the biological process of *niche construction*, whereby organisms modify their heritable environment. The concept of *niching* has been successfully used in EC previously, particularly in problems requiring multiple solutions (Mahfoud, 1995). However, niching in EC is primarily about maintaining stable sub-populations to improve the efficiency and efficacy of search – in general these methods do incorporate the biological concept of niche construction in their methodology, as is the case with the methods described in this paper. Before explaining the concept in more detail, we give a brief overview of the concept of a niche and niche construction.

Niches

In broad terms, biological environments have two main properties that determine the distribution and abundance of organisms: *conditions* and *resources*. Conditions are physiochemical features of the environment (e.g. temperature, pH, wind speed). An organism’s presence may change the conditions of its local environment (e.g. one species of plant may modify local light levels so that other species can be more successful). Conditions may vary in cyclic patterns or be subject to the uncertainty of prevailing environmental events. Conditions can also serve as stimuli for other organisms. Resources, on the other hand, are consumed by organisms in the course of their growth and reproduction. One organism may become or produce a resource for another through grazing, predation, parasitism or symbiosis, for example.

For any particular condition or resource, an organism may have a preferred value or set of values that favour its survival, growth and reproduction. Begon et al. (2006) define three characteristic curves, which show different “viability zones” for survival, growth and reproduction (Fig. 1).

The complete set of conditions and resources affecting an organism represent its *niche*, which can be conceptualised as a hypervolume in n -dimensional space. As an example, for two conditions c_1 and c_2 , two different types of species




Figure 2: Example exclusive and overlapping niche areas for a two-dimensional set of conditions.

relationships are shown in Fig. 2. The shaded area represents the viability zone for the species. A species will only survive if conditions are maintained within this shaded area. A relatively large distance in any single dimension denotes a generalist *in that dimension* (s_1 is relatively generalist in c_2), specialists have small distances (s_3 is more specialised in both c_1 and c_2). This size is referred to as *niche width*, and may vary for each dimension. If the mean viability zones overlap in a particular dimension, multiple species can co-exist within the range of overlap.

Competition and other species interactions are important in determining habitat distribution. Niche differentiation can permit coexistence of species within a biotope. Higher number of species can coexist by utilising resources in different ways. It is reasonably well understood in Biology how these mechanisms give rise to species diversity and specialisation.

The challenge addressed in this paper is to devise useful ways of employing these mechanisms in non-biological contexts. An important problem is in devising *appropriate* mappings between conditions and resources, and establish trade-offs for an individual's survival based on tolerances to specific conditions in order to enhance the quality and diversity of output in a creative generative system.

Niche Construction

Niche construction is the process whereby organisms change their own and each other's niches. They do this by modifying or influencing their local environment. Proponents of niche construction argue for its importance in understanding the feedback dynamics of evolutionary process in nature (Odling-Smee et al., 2003). By modifying their niche, either reinforcing or degrading it, organisms provide a heritable environment for their offspring. Hence niche construction can create forms of feedback that modify the dynamics of the evolutionary process, because ecological and genetic inheritance co-influence the evolutionary process. Computational models of niche construction show that it can influence the inertia and momentum of evolution and introduce or eliminate polymorphisms in different environments (Day

et al., 2003). Other models have demonstrated that a simple niche constructing ecosystem can support homeostasis and bi-stability similar to that of Lovelock's popular *Daisyworld* model (Dyke et al., 2007).

Whereas standard evolutionary algorithms tend to converge to a single (sub)-optimum, niche construction can promote diversity and heterogeneity in an otherwise fixed and homogeneous evolutionary system. In creative systems where the design of an explicit fitness function may be difficult or impossible, niche construction provides an alternate mechanism to explore a generative system's diversity over more traditional methods, such as the IGA. An "ecosystemic" approach to creative systems recognises that multiple designs may be equally valid and interesting, the emphasis shifting from single optimised solutions to the exploration of appropriate novelty offered through the feedback dynamics of an evolutionary ecosystem (McCormack, 2007).

Processes such as niche construction may serve as a type of "design pattern" (Gamma, 1995) that facilitates the building of creative evolutionary systems. To illustrate the utility of niche construction, we will describe a series of experiments where niche construction influences the structure and variation of the creative artefacts produced in an agent-based line drawing system.

Case 1: Line Drawing Agents

We will consider a simple creative system that autonomously draws lines with ink on a page. This system is inspired by Mauro Annunziato's *The Nagual Experiment* (Annunziato, 2002), which consisted of simple line drawing agents controlled by stochastic processes. In Annunziato's original system he changed the global characteristics of the drawings produced through manual adjustment of line-drawing probability parameters, such as fecundity, mortality and curvature. The resulting drawings have been acknowledged as artistically interesting and demonstrate the richness of creative output possible from a relatively simple generative specification.

Our system consists of a population of haploid line-drawing agents who inhabit a two-dimensional drawing surface or *canvas*. The canvas is initially blank (white). Agents roam over the surface, leaving a trail of black ink that marks out the path they travel. If a drawing agent intersects with an existing line, drawn either by itself or another agent, it dies. An agent may undergo reproduction during its lifetime, with offspring placed adjacent to the parent. The canvas is seeded with a small initial population of *founder agents*, initialised with uniformly distributed random genomes, that proceed to move, draw and reproduce. There is no limit to the number of offspring an agent may have, but in general the lifespan of agents decreases as the simulation progresses since the density of lines becomes greater, making it increasingly difficult to avoid intersection with existing lines. Eventually the entire population dies out (predominantly due to the intersec-

tion rule), and the image is finished. This finished drawing represents the “fossil record” of all the generations of lines that were able to live over the lifetime of the simulation.

In this first experiment, agents have no sensory information about their environment, for example they cannot detect proximity to an existing line or other agent. Thus, the characteristics of the line an agent draws are determined by genetics, with the genome serving as the control parameters of a stochastic process. An agent’s genome is specified by the following alleles, each represented as a normalised floating point value:

curvature (σ), controls the rate of curvature of the line ($\frac{d\theta}{dt}$, where θ is the heading direction). Curvature varies from a straight line (0) to a maximum curvature rate (1);

irrationality (r), controls the rate and degree of change in the rate of curvature according to a stochastic algorithm (detailed below, see also Fig. 3);

fecundity (f), the probability of the agent reproducing at any time step. New agents are spawned as branches from the parent;

mortality (m), the probability of the agent dying at any time step;

offset (ϕ), the offset angle of child filaments from the parent;




Figure 3: Individual line drawing agents with different measures of irrationality. Note that the ‘die if intersect’ rule has been turned off for these examples.

In addition each agent maintains *state* information which includes the current position on the canvas, heading direction, speed and current rate of curvature. Changes to the rate of curvature are determined by the **curvature** and **irrationality** alleles, with the overall rate of change given by

$$\frac{d\theta}{dt} = \sigma + \text{fracSum}(\mathbf{p}, k \cdot r)^{0.89r^2}, \quad (1)$$

where \mathbf{p} is the agent’s current position, k a constant known as the *octave factor*, and fracSum a function that sums octaves of Perlin (2002) 2D noise. This function was chosen as it gives band limited, continuous stochastic variation with

second order continuity, and is statistically invariant under affine transformation. Increasing r (irrationality) increases the octaves of noise, changing the rate of change in direction in increasingly finer detail. Fig. 3 shows the effects of varying the irrationality allele, r , over its normalised range.

This system was run a number of times varying the random number seed and location of founder agents on the blank canvas. At each time step the **fecundity** and **mortality** alleles determine probabilistically if an agent will die or reproduce. In the case of reproduction, child agents are placed next to the parent line, with their heading determined by the offset allele (ϕ). A child agent’s genome may undergo mutation (modification of an allele by adding a Normally distributed random number with mean 0). Additionally, children have a short gestation period before they begin to draw, allowing the parent to continue drawing past the point where reproduction took place, avoiding intersection with their offspring.

The images that emerge from this process demonstrate a wide variety of output possible from this system (two sample images are shown in Fig. 4). While there is no explicit fitness function or evaluation, implicit agent fitness is determined by a combination of genetics and environment. Importantly, the environment is constantly changing. As drawing progresses, it becomes increasingly difficult to reproduce and live, since the probability of intersecting with an existing line typically becomes higher as more lines crowd the canvas.

While the images produced by this system are interesting, in general they lack a changing dynamic or visual counterpoint, that is, they are largely homogeneous in structure, or have progressive changes that take place as genes mutate through drift. Much of the overall structure is determined by the founder lines, who can carve up large areas of blank canvas for themselves and their offspring, preventing other lines from entering. Genetically similar offspring continue to reproduce inside these boundaries until the space is filled.

Case 2: Line Drawing with Niche Construction

In a second experiment we tested the hypothesis that by introducing an ecosystemic process of niche construction into the system, the overall diversity and heterogeneity of images produced by the system could be significantly enhanced. To do this, each agent was given an additional allele in its genome: a local density preference δ_i (a normalised floating point value). This defines the agent’s preference for the density of lines already drawn on the canvas in the immediate area of its current position, i.e. its *niche* (Fig. 5). In a preferred niche, an agent is more likely to give birth to offspring and has a better chance of survival. As children inherit their parent’s genes they are more likely to survive as they have a similar density preference. So in a sense, parents may construct a niche and pass on a heritable environment well-suited to their offspring.



Figure 4: Two sample outputs from the line drawing system (without niche construction).




Figure 5: The niche construction mechanism for drawing agents, who try to construct a niche of local density that satisfies their genetic preference.

For each agent, i , δ_i defines its preferred niche. Local density, defined as the ratio of inked to blank canvas per unit area, is measured over a small area surrounding the agent at each time step. Proximity to the preferred niche determines the probability of reproduction, given by

$$Pr(rep) = f_i \cdot \cos^\omega(\text{clip}(2\pi(\Delta p_i - \delta_i)), -\frac{\pi}{2}, \frac{\pi}{2}), \quad (2)$$

where Δp_i is the local density around the point p_i , the agent's position, ω a global parameter that varies the effective niche width, f_i is the agent's fecundity and clip is a function that limits the first argument to the range specified by the next two. Being in a non-preferred niche similarly increases the probability of death.

Founder agents begin with a low density preference, uni-

formly distributed over $[0, 0.2]$. Beginning the drawing on a blank canvas means that only those agents who prefer a low density niche will survive. As the drawing progresses however, more ink is added to the canvas and agents who prefer higher densities will prosper. As with the previous experiment, at birth the agent's genome is subject to the possibility of mutation (proportional to the inverse of the genome length), allowing offspring to adapt their density preference and drawing style as the drawing progresses. Eventually the population becomes extinct, since higher density favouring agents don't have much room to move, and the drawing finishes. Some example drawings are shown in Fig. 6. Notice the greater stylistic variation and heterogeneity over the images shown in Fig. 4.

Analysis and Discussion

Visually, the examples appear to show that by adding niche construction, the line drawing system is capable of producing images with greater heterogeneity, variation in density, counterpoint and overall visual interest (Fig. 7). We might even be tempted to say it is more creative.

To support this intuition, a number of images produced using the niche constructing and non-niche constructing versions were analysed statistically. A total of 40 images were sampled: 20 niche constructed and 20 non-niche constructed. For each image, the mean density ($\bar{\Delta}$) and variance of density over the entire image was computed. Then for each set (non-niche constructed, niche constructed) the variance of mean density and the mean density variance was calculated. Table 1 summarises this analysis. p -values were calculated using a Welch t-test. As shown in the table, niche




Figure 6: Two sample outputs from the line drawing system with niche construction.



Figure 7: Detail from two drawings, showing density variation (left) without niche construction, and (right) with niche construction.

constructed images exhibit a far greater variation in overall density (by a factor of 3.83). Significantly, the density variation *over each image* is, on average, 4.31 times greater for the niche constructed over non-niche constructed drawings.

	Non NC	NC	<i>p</i> -value
Number of Images	20	20	–
Variance of $\bar{\Delta}$	0.00298	0.0114	0.0634
Mean Variance	0.0140	0.0604	1.57×10^{-10}

Table 1: Density variation between non-niche constructed and niche constructed drawings.

Analysis of the mean agent density preference, $\bar{\delta} = \frac{1}{n} \sum_{i=1}^n \delta_i$, at each epoch shows an overall adaption to the mean image density ($\bar{\Delta}$) over the lifetime of the drawing,

indicating that agents evolve to fit niches (Fig. 8). On average, agents favour slightly denser niches than currently exist (the line in the figure is always positive), we infer this is because an agent’s density measure is always centred around the agent’s current location, and this will necessarily include parts of the images with lines drawn (even if only the agent’s own trail). The value of $\bar{\Delta}$ tends to increase over the life of the drawing. This is not surprising, as there is no mechanism for an agent to *reduce* the density of its niche². The best any parent can do is carve out the largest possible border around empty space, so that its offspring can grow without fear of intersecting with other parents or their offspring.

Conclusions and Future Work

We have demonstrated how the ecological “design pattern” of niche construction can be used to enhance the creative output of a generative line-drawing system. Elsewhere, (McCormack and Bown, 2009), we have also applied a similar process in the sound domain, leading to on-going change in an agent-based sound generation system. While it may be premature to suggest the generality of this method, our on-going experiments demonstrate that with the appropriate design, niche construction can introduce heterogeneity and useful variation into creative generative systems.

The line-drawing agents described in this paper have only one way to sense their environment: through their density preference. A more sophisticated system might give agents greater sensory capabilities so that they can better optimise

²An observed (short-lived) strategy is to draw a closed circular area and not place any offspring in it, but this only generates a low-density niche after death!

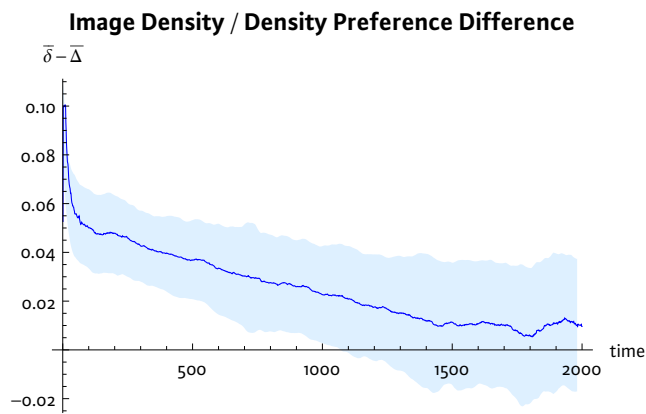


Figure 8: Difference between mean image density and mean agent density preference averaged over 40 runs. The standard deviation is shown in light blue.

their niche construction to their environment. For example being able to sense proximity to another line would allow more graphically complex strategies to evolve.

Additionally, the agents are limited in their productive utilisation of evolution, as any adaptation must take place over the life of a single drawing. Typically, $10^3 - 10^5$ offspring may be produced in a single image, but less than 10 – 30 generations from the initial parent. Essentially, all lines are of the same species. An improved strategy would be to allow different species of line-drawing agents to be pre-evolved on test canvases, permitting better optimisation for different density niches and inter-species interactions. These pre-evolved species could then share a common drawing canvas in order to produce a more complex finished drawing, better adapted to their specific niche requirements. We are currently exploring this idea. One can imagine that the next generation of artist's drawing systems could incorporate such pre-evolved drawing agents as "intelligent brushes"; the artist selecting from a palette of pre-evolved styles and applying them to the canvas at various stages. Agents with different niche density preferences try to draw in order to construct their preferred niche, but their interactions with each other could result in the emergence of competitive or cooperative strategies.

In summary, we believe that niche construction is a useful technique that can be successfully exploited in generative creative systems to enhance the dynamics and heterogeneity of output produced. The ecosystemic approach favoured in this paper is in contrast to previous IGA or fitness-based GA systems aimed at search or optimisation to singular outcomes or subjective criteria. The complex dynamics of ecosystem processes are a source of rich and varied inspiration that has much to offer as we develop autonomous creative systems.

Acknowledgements

This work was supported by Australian Research Council Discovery Projects grant DP0877320.

References

- Annunziato, M. (2002). The nagual experiment. <http://www.plancton.com/papers/nagual.pdf> (Accessed 11 November 2008)
- Baker, E. and Seltzer, M. I. (1994). Evolving line drawings. In *Graphics Interface '94*, pages 91–100, Banff, Canada.
- Basalla, G. (1998). *The Evolution of Technology*. Cambridge University Press, Cambridge, MA.
- Begon, M., Townsend, C., and Harper, J. (2006). *Ecology: from individuals to ecosystems*. Wiley-Blackwell.
- Bird, J., Husbands, P., Perris, M., Bigge, B., and Brown, P. (2008). Implicit fitness functions for evolving a drawing robot. In Giacobini, M., Brabazon, A., Cagnoni, S., Caro, G. D., Drechsler, R., Ekárt, A., Esparcia-Alcázar, A., Farooq, M., Fink, A., McCormack, J., O'Neill, M., Romero, J., Rothlauf, F., Squillero, G., Uyar, S., and Yang, S., editors, *EvoWorkshops*, volume 4974 of *Lecture Notes in Computer Science*, pages 473–478. Springer.
- Birkhoff, G. D. (1933). *Aesthetic Measure*. Harvard University Press, Cambridge, MA.
- Boden, M. A. (2004). *The creative mind: myths & mechanisms*. (2nd edition), Routledge, Oxon.
- Carroll, S. (2005). *Endless forms most beautiful: The new science of evo devo and the making of the animal kingdom*. WW Norton & Company.
- Day, R. L., Laland, K. N., and Odling-Smee, J. (2003). Rethinking adaptation: the niche-construction perspective. *Perspectives in Biology and Medicine*, 46(1):80–95.
- Dyke, J. G., McDonald-Gibson, J., Di Paolo, E., and Harvey, I. R. (2007). Increasing complexity increases stability in a self-regulating ecosystem. In *Proceedings of IXth European Conference on Artificial Life, ECAL 2007*, LNCS 4648, pages 133–142.
- Gamma, E., Helm, R., Johnson, R., and Vlissides, J. (1995). *Design patterns: elements of reusable object-oriented software*. Addison-Wesley professional computing series. Addison-Wesley, Reading, Mass.
- Gould, S. J. (2002). *The Structure of Evolutionary Theory*. Harvard University Press, Cambridge, MA.
- Griffiths, T. D. (2008). Capturing creativity. *Brain*, 131(1):6–7.
- Laland, K. N., Odling-Smee, J., and Feldman, M. W. (1999). Niche construction, biological evolution and cultural change. *Behavioral and brain sciences*, 23(01):131–146.
- Machado, P., Romero, J., and Manaris, B. (2008). Experiments in computational aesthetics. In Romero, J. and Machado, P., editors, *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, Natural Computing Series, pages 381–415. Springer.

- Mahfoud, S. W. (1995). A comparison of parallel and sequential niching methods. In *Proceedings of the 6th International Conference on Genetic Algorithms*, page 143, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Maynard Smith, J. and Szathmáry, E. (1995). *The major transitions in evolution*. W.H. Freeman Spektrum, Oxford; New York.
- McCormack, J. (2005). Open problems in evolutionary music and art. In Rothlauf, F., Branke, J., Cagnoni, S., Corne, D. W., Drechsler, R., Jin, Y., Machado, P., Marchiori, E., Romero, J., Smith, G. D., and Squillero, G., editors, *EvoWorkshops*, volume 3449 of *Lecture Notes in Computer Science*, pages 428–436. Springer.
- McCormack, J. (2007). Creative ecosystems. In Cardoso, A. and Wiggins, G., editors, *Proceedings of the 4th International Joint Workshop on Computational Creativity*, pages 129–136.
- McCormack, J. (2010). Creative evolution. In *preparation*.
- McCormack, J. and Bown, O. (2009). Life’s what you make: Niche construction and evolutionary art. In Giacobini, M., Brabazon, A., Cagnoni, S., Caro, G. A. D., Ekárt, A., Esparcia-Alcázar, A., Farooq, M., Fink, A., Machado, P., McCormack, J., O’Neill, M., Neri, F., Preuss, M., Rothlauf, F., Tarantino, E., and Yang, S., editors, *EvoWorkshops*, volume 5484 of *Lecture Notes in Computer Science*, pages 528–537. Springer.
- Nowak, M. A. (2006). *Evolutionary Dynamics: exploring the equations of life*. The Bekknap Press of Harvard University Press, Cambridge, Massachusetts, and London, England.
- Odling-Smee, J., Laland, K. N., and W., F. M. (2003). *Niche Construction: The Neglected Process in Evolution*, Monographs in Population Biology. Princeton University Press.
- Perlin, K. (2002). Improving noise. *ACM Transactions on Graphics (TOG)*, 21(3):681–682.
- Perlovsky, L. I. (2010). Intersections of mathematical, cognitive, and aesthetic theories of mind. *Psychology of Aesthetics, Creativity and the Arts*, 4(1):11–17.
- Ramachandran, V. S. (2003). *The Emerging Mind*. Reith lectures; 2003. BBC in association with Profile Books, London.
- Ramachandran, V. S. and Hirstein, W. (1999). The science of art: A neurological theory of aesthetic experience. *Journal of Consciousness Studies*, 6:15–51.
- Stauderk, T. (2002). *Exact Aesthetics. Object and Scene to Message*. PhD thesis, Faculty of Informatics, Masaryk University of Brno.
- Sternberg, R. J. (1999). *Handbook of creativity*. Cambridge University Press, Cambridge, U.K.; New York.
- Svangård, N. and Nordin, P. (2004). Automated aesthetic selection of evolutionary art by distance based classification of genomes and phenomes using the universal similarity metric. In Raidl, G. R., Cagnoni, S., Branke, J., Corne, D., Drechsler, R., Jin, Y., Johnson, C. G., Machado, P., Marchiori, E., Rothlauf, F., Smith, G. D., and Squillero, G., editors, *EvoWorkshops 2004*, volume 3005 of *Lecture Notes in Computer Science*, pages 447–456. Springer.
- Takagi, H. (2001). Interactive evolutionary computation: Fusion of the capabilities of ec optimization and human evaluation. *Proceedings of the IEEE*, 89:1275–1296.