

Ecosystemic methods for creative domains: niche construction and boundary formation

Oliver Bown

Centre for Electronic Media Art
Monash University
Melbourne, Australia, 3800
Email: ollie@icarus.nu

Jon McCormack

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

Taras Kowaliw

Institut des Systèmes Complexes -
Paris Île-de-France
Centre national de la recherche scientifique,
57-59 rue Lhomond, 75005, Paris, France

Abstract—This paper examines how models from ecology and evolutionary multi-agent systems can be used for creative applications. We outline a minimal, general specification for an ecosystem to be applied to the production of creative artefacts, based on spaces, materials, features and actions. Using this specification, we consider how basic evolutionary principles based on niche construction can be used to establish the emergence of heterogeneous structures in the simplest imaginable cases to which our framework applies. We investigate the minimal conditions for the formation and maintenance of boundaries, and then apply the results of our initial models to the design of a creative artwork, contrasting differences between niche-constructing and neutral evolutionary models.

I. INTRODUCTION

The adoption of generative methods in the creative arts has grown consistently over the past few decades primarily, but by no means exclusively, through the use of digital technologies. A natural continuation of this trajectory is the application of multi-agent systems to traditional creative domains using populations of agents that adapt and self-organise to produce novel structures. Multi-agent models, and more specifically models based on biological ecosystems which combine evolution and niche construction, offer a compelling new approach to the generation of novel forms in art, design and music. Although this idea is attractive, little research has been done to suggest how such an approach might be formulated and put into practice.

This paper investigates how ideas from ecosystem modelling could be employed in the automated emergence and maintenance of *structure*. One minimalist interpretation of structure requires the presence and maintenance of a boundary between different environmental properties. We consider the emergence of heterogeneous populations in the simplest relevant scenario, with boundary divisions in a single dimension. The results are discussed in terms of how the structuring processes can be extended into higher dimensions, and we present results of experiments in 2D discrete environments based on a single environmental resource. These principles are then extended to a more complex, two species engineering web that provides the basis for introducing novelty into a generative drawing system using an agent-based ecosystem model.

II. ECOSYSTEM MODELLING AND NICHE CONSTRUCTION

Over the last few decades, a number of researchers have claimed that the causal interactions between species and their environments is bidirectional or cyclic, in opposition to the traditional unidirectional view of organism adapting to environments [1]. The concept of *Niche Construction*, introduced by Lewontin, and championed by biologists such as Odling-Smee, Laland and Feldman, describes the process whereby organisms create and modify their own and each other's niches. This, proponents argue, allows the establishment of a heritable environment for their offspring. By changing their environment, organisms can positively or negatively affect their own survival and reproduction opportunities (and possibly that of other species). In this way, niche construction is regarded as an initiator of evolutionary change, where both ecological and genetic inheritance influence selection.

Largely independent of evolutionary considerations, *ecosystem modelling and simulation* focuses on the interaction between species and their environment. Ecosystems are often modelled without incorporation of genetic or evolutionary processes, as ecologists are more concerned with modelling over time scales where evolution has little or no effect. Here, ecologists use the term *ecosystem engineering* [2]: essentially niche construction in an ecosystem context.

Artificial Life research has often focused on combining simple multi-agent ecosystem and evolutionary models, for purposes such as the study of evolutionary dynamics [3], niche proliferation and inter-species interaction [4], emergent dynamics [5] or complex human-computer interactions in an artistic context [6], [7]. The use of individual-based models in ecology and evolution are reviewed in [8].

Several studies of niche construction in the Artificial Life literature have focused on spatial re-implementation of existing models of evolutionary interaction, and specifically niche construction (e.g. [9], [10], [11], and the earlier work of Axelrod [12]). These spatial models confirm the generality of the central implications of niche construction, and also the impact of spatialisation, which introduces additional self-organisation as populations of agents strive to inhabit suitable niches under topological constraints. For the purpose of this paper we therefore distinguish three interdependent processes: evolutionary

adaptation, niche construction and self-organisation.

III. MULTI-AGENT APPROACHES TO GENERATIVE ART

Generative art involves the creation and use of automated methods (usually computer programs) to autonomously generate artworks, shifting the artist's role to that of meta-creator [13]. Typically, a generative art system can produce a very large set of outputs which will have a stylistic commonality expressed by the generative mechanism that created them. One challenge for generative art is to find methods that simultaneously maximise both the quality and diversity of the output produced by a given system, increasing the creative autonomy of the system. This is a difficult problem since both value and novelty – the central, co-dependent features of creative output – are hard to capture in a single formal model. Another challenge is to find representations and methods that enhance the power of the human artist, or community of artists, to creatively innovate using generative tools, which is where our interest lies in this paper.

A multi-agent approach to generative art and music allows for the complex pattern-generating potential of collective self-organisation. This approach is characterised by an ability to operate across multiple scales, generating pattern in a bottom-up manner from local structure. Additionally, it allows the mixing of different strategies in a shared common environment. At the same time however, this approach is inherently more complex as it requires structure to emerge from the interaction of lower-level elements, making it difficult to directly define how a system will act at the global level (this is of course the original source of interest in such systems for many researchers).

Reynolds' *boids* flocking algorithm [14] is recognised as an early creative example of a multi-agent system directly inspired by nature. The algorithm was adapted to a musical domain by Blackwell and Young, who draw an analogy between the self-organising behaviour of a swarm or flock and an improvising ensemble [15]. In other examples, a more abstract use is made of the principle of self-organising in multi-agent systems, such as the distributed drawing systems studied by Driessens and Verstappen [16], Eldridge and Dorin [17], McCormack [18], and Greenfield [19]. Dorin provides an overview of a variety of ecosystem models in the arts [20], covering a number of more conceptual applications.

Evolution provides a method for generative systems to adapt to external requirements. This has been convincingly evidenced in architectural design where the system's requirements amount to a traditional evolutionary optimisation problem and the design challenge involves searching a space of acceptable solutions applying human design decisions [21]. As in nature, the solutions discovered by evolutionary optimisation of complex multi-constraint problems can be unexpected (e.g. [22]) illustrating the power of evolution to discover designs outside the scope of conventional design thinking, maximising the tenet that function defines form.

In visual art and music, however, specifying global requirements is a more arduous task: it has proven immensely difficult

to formalise aesthetic preferences, and the kinds of high-level aesthetic measures of artworks or musical works that have been proposed so far are either underspecified (allowing too many unacceptable outcomes), over specified (do not allow a sufficient degree of variety) or both (allow variation but along unsatisfactory dimensions). A reasonably successful alternative approach, however, is the interactive genetic algorithm (IGA), which couples a user's preference (instead of an aesthetic measure) to the evolution of a generative process.

In [23] we argue that evolution driven by changes from within a system, such as those established through coevolution or the coupling of niche construction and adaptation, could also lead to elegant and aesthetically valuable outcomes in the hands of an experienced creative practitioner, but that this still requires the ability for the artist to manipulate the system effectively. The following section proposes an approach to organising creative domains so that they could be subject to ecosystem modelling methods, with this flexibility in mind.

IV. ECOSYSTEM MODELS AS CREATIVE TOOLS

A. Feature Spaces as Niches

A *niche* is the set of *conditions* (physiochemical environmental features) and *resources* (biotic material consumed or processed by an organism) within which an organism makes its habitat. Species may have a finite range of acceptable tolerances for conditions and resources, which define a multi-dimensional hypervolume of viability. Niche construction is the process whereby organisms change their own and each other's niches via environmental modification. Computational models of niche construction show that it can influence the inertia and momentum of evolution and introduce or eliminate polymorphisms in different environments [24].

To create niche-constructing ecosystem models, agents typically express a genetically encoded preference for environmental quantities. An additional genetic trait determines how the agent modifies their local environment, typically to manipulate favourable environmental values (although it is quite possible that environmental change may impact negatively, e.g. excessive waste production).

In a creative context, environmental conditions and resources, along with the methods agents have of manipulating them, must be established by the system's designer. For aesthetic artefacts, the environment is typically the artefact itself, whose structure can be described by sets of local features, typically corresponding to models of human perception. Features can occur at different scales corresponding to high and low-level perceptual phenomena. In music, for example, an audio signal can be interpreted spectrally at a fine granularity, the spectral data can be used to find sound event onsets, and these onsets can be used to infer the tempo of the music. As localised, complex mappings of a structure into a perceptually meaningful domain, we believe that such features provide the perfect bridge between a creative users' aesthetic interest and the environmental conditions from which an agent's niche might be determined, mirroring the coupling between preferences and outcomes found in IGAs.

B. Ecosystemic Specification of a Creative Domain

On this basis we propose the following set of elements to complete a minimal ecosystemic specification of any creative domain:

1) *A space*: A creative domain consists of a topologically constrained *space*. For example, this may be the one dimensional timeline of a musical work, the two dimensional pixel grid of an image, or the three-dimensional space of a sculpture. The space can be described by a set of locations (often discrete) and a set of neighbourhood proximity relations between those locations, including the boundary conditions. The space, its topology and neighbourhood relations, determine the *environment* in which the ecosystem exists.

2) *Materials*: An environment consists of materials. In a 2D image, the materials might be pixels, but they could also be shape primitives, if using a vector-based format. In a musical context, they could be the individual samples that make up a waveform, or other generative objects such as sine wave grains or sound files that are mixed together additively. In each case the materials are bound to locations, such that a set of locations and the materials they contain describe an artwork. A 2D image may consist of a 2D grid of cells each containing an RGB value, and a musical work could consist of a set of discrete time steps each containing zero or more MIDI¹ events. Many other possibilities are available, associated with different existing practices (and associated representation formats) across the digital arts. The space and materials together define an art object.

3) *Features*: Each location in an art object can be analysed using feature extraction, as discussed above. In a musical work, the MIDI notes that make up the material of the artwork ultimately generate sounds which can produce different levels of consonance, different spectral qualities, or temporal phenomena such as event onsets or tempi. In a 2D image, combinations of pixels might produce high-level phenomena such as coherent gradients, edges or shapes. Note that these conditions can be derived directly from the materials in the environment. Thus given a set of feature extractors, one can specify a relationship between a concrete structure (the topology of an environment and its material contents) and a derived interpretation of that structure. We present these as mappings from an environment (space and materials) to a space of living conditions from which an organism's niche can be defined. For example, an agent occupying a location in a 2D image might be adapted to an edge niche, which would be satisfied if the pixels surrounding that location were such that they formed an edge. This could be analogous to specific habitats provided by the branches of a tree, or the edge of a river, which can be thought of as exhibiting qualities that are naturally derived from the materials that make them up.

4) *Actions*: For agents to be able to make modifications to their environment they must act upon the materials that make up their environment at or around their location. By

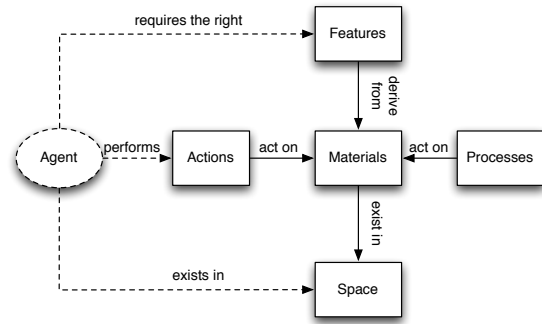


Fig. 1. Relations between elements in ecosystemic specification of a creative domain.

doing so they will modify the features at that location, thus affecting the environmental conditions that make up their niche. The modifications an agent can make is defined by a set of possible actions associated with the definition of a space and its materials. Actions operate on materials, never directly on features. Actions may also be constrained in various ways. For example, adding an element at one location may require moving it from another location, if a conservation rule is being applied. Equally in the absence of a conservation rule, it should be possible to create or remove elements. Actions may incur a cost to the agent or may be constrained in their efficacy.

5) *Processes*: Finally, natural processes may occur in environments, modifying the materials in those environments in what may be understood as an abiotic way, e.g., corresponding to decay, energy input from the sun, natural rhythms such as the alternating conditions of day and night, or the diffusion of chemicals through an environment. Likewise, an artificial image can be subject to functionally similar processes, causing the image to change structure over time. An artificial musical environment can be subject to chains of effects processors prior to the extraction of features. In the first case this corresponds to a dynamic process, meaning that the conditions are in constant dynamical flux, whilst in the second case this corresponds to a more complex instantaneous mapping from materials to features (via a temporary intermediate material or feature), possibly involving action at a distance. For example, a reverb or delay effect would mean that events occurring earlier in the timeline could influence features extracted at locations later in the timeline.

The set of relations between these elements is presented in Figure 1. Following this specification, any system can then be subject to a creative ‘ecosystemification’ process. At this stage, we have not specified agent behaviour, or any relationship between niches and fitness. For the time being, we will seek generic approaches to the design and configuration of agent behaviour, niche preference and fitness, species interactions and intra-species population dynamics. Figure 1 indicates the basic ways in which generic agents would interact with such a system.

¹Musical Instrument Digital Interface, a standard protocol for musical note and control information.

V. STUDIES

Our studies investigate multi-agent models of niche construction and evolution as a means to establish heterogeneity, local specialisation and structure in art objects. We test the conditions under which quasi-stable environmental heterogeneity can be formed through evolutionary niche construction, being maintained by two distinct populations operating on the same environmental resource, and in some cases emerging from a homogeneous population and environment. To do this we demonstrate processes of boundary formation in simple 1D and 2D greyscale images using the simplest combination of the elements we have defined in Section IV-B. These studies explore the tolerances of boundary formation and stability in these simple cases. We then apply these ideas to the production of a visual artwork, demonstrating the effects of niche construction-based heterogeneity.

A. 1D Greyscale Model

We begin with a bounded 1D environment of 100 discrete connected cells, containing a single material: a normalised brightness level, v , over the interval $[0, 1]$. We consider only a single feature: the mean brightness value (\bar{v}) across a window with width $(2w + 1)$ centred on the location at which the feature is being measured (note the difference between this feature which is the mean brightness, and the material which is the exact brightness at a given location). The environment is populated by a single agent at each cell. Each agent contains a real-valued gene encoding preference for a specific brightness level. The fitness, f_i of the agent at location i is determined by their adaptedness to the average brightness feature at their location, given by:

$$f_i = e^{-10|p_i - \bar{v}_i|} \quad (1)$$

where p_i is the feature preference and \bar{v}_i is the feature value at the agent's location.

At each time-step, n niche construction events occur followed by one tournament event. At each niche construction event a random agent is chosen. It chooses a random cell within w cells of its location (near the boundary a selection is made only from the available cells) and causes a random change in the brightness value at that location, drawn from a Gaussian distribution with standard deviation of 0.01. The agent then accepts this change if it results in a more preferable health, and rejects it otherwise. At each tournament event, two agents are chosen at random and the less fit agent is replaced by an offspring from the more fit agent. The winner is chosen at random if the agents are equally fit. Offspring have a probability of 0.1 of having a creep mutated preference, drawn from a Gaussian distribution with standard deviation 0.01. Selection of agents for tournaments are either chosen entirely at random (global tournament selection), or alternatively one agent is chosen at random, and a random neighbour of that agent (any agent within w cells) is then chosen (local tournament selection).

Different values of n (1, 2 and 10) and w (1, 2 and 4) were compared, as well as local versus global tournament

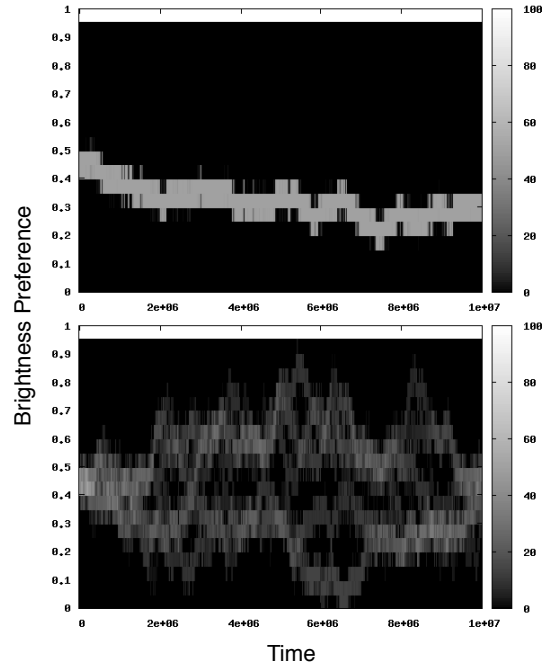


Fig. 2. Histograms showing the distribution of agents' preferred brightness over time, in the case of global tournament selection (top) and local tournament selection (bottom). In both cases $w = 2$, $n = 10$.

selection, random (uniformly distributed in the range $[0,1]$) versus constant (fixed at 0.5) preference initialisation, and random versus constant brightness value initialisation. All simulations were run for ten million time steps.

The results show that, as might be expected, global tournament selection produces strongly homogeneous populations and environments. The preference of the population drifts randomly (and the environment changes with it) but maintains a low degree of variation. Local tournament selection, on the other hand, results in the emergence of stable boundaries separating heterogeneous groups of genetically similar agents, each of which is adapted to its environment. This distinction was observed across variation in all the other settings mentioned, although the heterogeneity of the local tournament selection models diminished with increasing n , and the rigidity of boundary positions over time diminished with increasing w . Figure 2 shows a comparison of histograms of agent preference for typical runs as they evolve over time for global (top) and local (bottom) tournament selection, with $w = 2$, $n = 10$. Figure 3 shows two typical evolutions of agent preferences in a local tournament selection model, beginning with a uniform population and environment, with $w = 1$, $n = 1$ (top) and $w = 2$, $n = 10$ (bottom). This shows each sub-population's preference (and environment) drifting randomly, but that the boundaries remain stable over time (more so in the first case).

The stability of boundaries can be understood by considering the potential for the offspring of an agent from one side of a boundary to successfully invade the other side. Agents of one group that are within w cells of the boundary cannot be

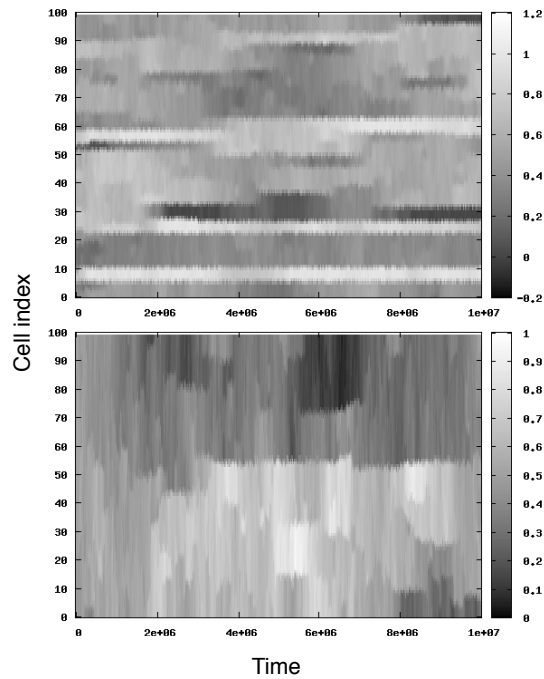


Fig. 3. Examples of heterogeneous populations emerging from neutral homogeneous environments with $w = 1$, $n = 1$ (top) and $w = 2$, $n = 10$ (bottom). The graphs show agent preference (shading) at each location (y-axis) over time (x-axis).

perfectly adapted to their environment, and the closer agents are to the boundary, the less fit they are. This means it is possible for an agent closest to the boundary to be defeated by an agent from the opposite side. However, the resulting child will be in an environment that is slightly less suitable, and will not be able to bring the environment around to its preference (as long as its preference hasn't mutated very close to the boundary value). The new agent will soon be replaced by a competitor from the other side. Thus increasing n (the relative degree of niche construction) increases the chances of an invader having some luck in transforming the environment and thus diminishes the degree of heterogeneity in the population.

This understanding of boundary stability was investigated by running versions of the local tournament selection model from a start state consisting of darker-preferring agents in one half of the environment, and lighter-preferring agents in the other half. The environment was initialised such that it matched this distribution of agent preferences. The degree of initial deviation either side of the neutral value (0.5) was varied, using values 0.01, 0.1 and 0.5. The results show that a greater gap between environments led to a stronger boundary, supporting the above description of boundary stability (Figure 4). In the case of the smallest deviation 0.01, it was possible for invaders with a mutation towards a more intermediate preference to take hold at the boundary (Figure 4, top).

As mentioned, boundaries between species also consistently emerged in the local tournament selection case. Figure 5 shows an instance of boundary emergence from a homogenous species/environment, indicating the growth of subpopulations

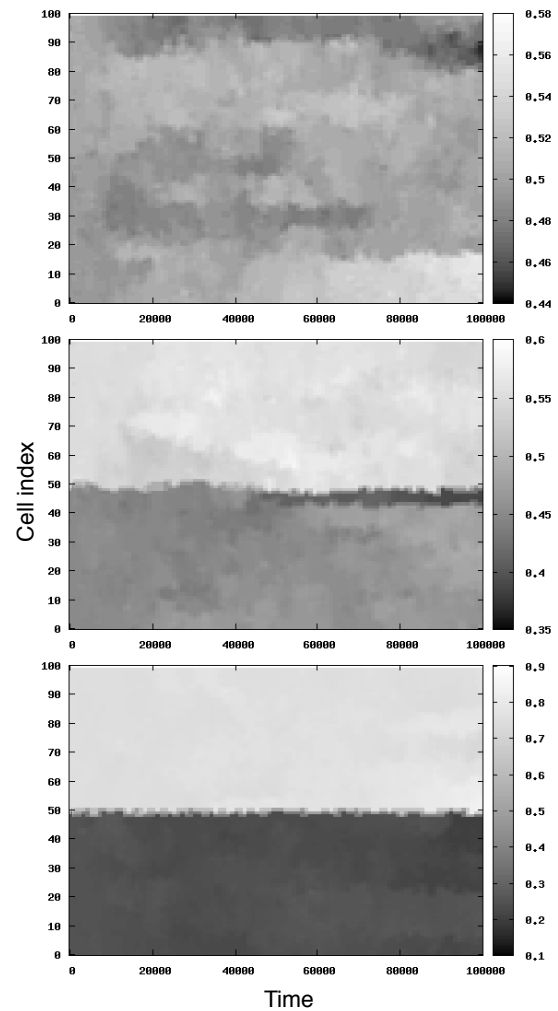


Fig. 4. Stability of boundary between two species with initial deviations from the neutral value of 0.01 (top), 0.1 (middle) and 0.5 (bottom), using $w = 2$, $n = 1$, local tournament selection. Note that the greyscale ranges are different in each graph, to accentuate the detail in each case.

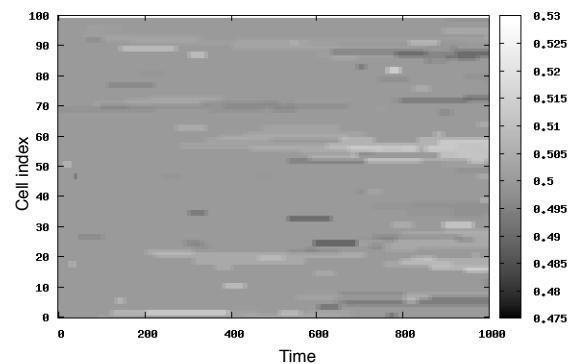


Fig. 5. Demonstration of early boundary emergence. Random mutations are originally established by chance in the absence of strong selective pressure during local tournament selection. As populations drift, boundaries become established. The figure is a close-up of the first 1000 time-steps of Figure 3 (bottom).

through very gradual drift. Whilst the drifting subpopulations remain similar to each other, no group has sufficient advantage

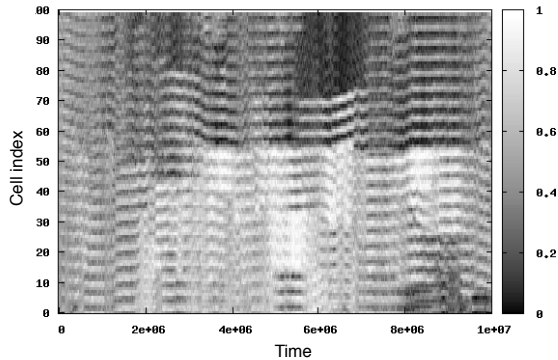


Fig. 6. The same simulation as Figure 3 (bottom) but representing the state of the environment rather than the preference of agents. The striped pattern is a result of agents collectively attempting to modify their local environments to produce a consistent average value.

over the other, and some simply grow by chance, leading to the emergence of a boundary which is then reinforced by the process described above.

Finally, as a side effect of the simple implementation of this model, it was noted that environmental niche construction by a homogeneous population of agents could lead to a heterogeneous patterning of the environment (Figure 6). Since an agent’s preference is for an average brightness value over a range specified by w , there are many ways that this environmental preference can be satisfied. One way would be for each of the locations to have exactly the desired brightness value, but it is equally possible for darker cells on one side to be balanced with lighter cells on the other side to produce the desired average value. In this way, alternating dark and light bands can satisfy the preferences of an entire homogeneous population. This self-organising alternating pattern was observed to occur frequently. It is recognised that this is an artefact of the simple and relatively arbitrary decision to average evenly over a feature window. If the average had been weighted by distance from the source cell then this alternating effect could not have emerged. However, similar patterning effects were also observed in a wide range of experiments not reported here. Although this is not necessarily a phenomenon that occurs in nature, it suggests a way that populations could generate interesting rhythmic patterning. It also implies that even in simple cases, a non-trivial relationship between preference and niche-construction behaviour emerges, which drives a process of environmental modification.

B. 2D Greyscale Model

The neighbourhood topology of one dimensional and two dimensional environments is qualitatively different, so the generality of the result was tested with a two dimensional model. In the case of a linear boundary between two distinct environments the results cited above apply equally well in two dimensions. However, the explanation of boundary stability suggests that an imbalance of any kind will result in instability. This is confirmed when we consider a start state consisting of two different environmental preferences as above, but configured such that one forms a circle enclosed by the

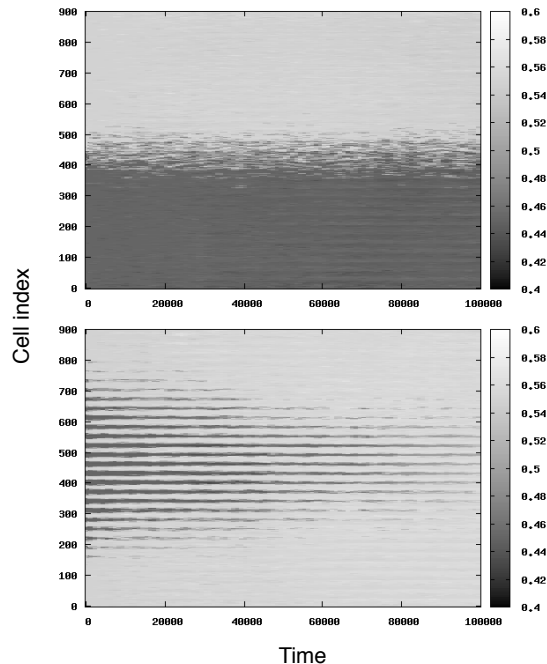


Fig. 7. Demonstration of the greater stability of a linear boundary as compared to a circular boundary. The 2D environment has been unfolded to a row major linear ordering and is shown evolving over time as in the previous figures. Both simulations used local tournament selection on a 30x30 grid with $w = 4$, $n = 10$ and preferences of 0.4 and 0.6 for the initial species (deviation of 0.1 from the neutral value).

other. Unlike a linear boundary where all things are equal, a population inhabiting a circle can be more easily invaded by the population surrounding it. At any location on the boundary, the number of neighbours belonging to the outer population is slightly greater than the number of neighbours belonging to the enclosed population, and the outer population will therefore be able to gradually drive the net effect of the combined niche construction towards their preferred niche. Figure 7 compares the evolution of environmental preferences in the linear and circular cases where all other settings are the same.

VI. ARTISTIC EXPERIMENTS

As the methods described in this paper are intended for enhancement of creative possibilities in generative artistic systems, we developed a number of simple artistic sketches that employ niche construction and environment modification in order to judge their utility. In a previous paper, we demonstrated how niche construction could be used to modify the behaviour of line drawing agents on a canvas [18]. By adding a genetically encoded preference for local image density, the agents were able to modify their behaviour to construct niches that suited their density preference and that of their offspring. The resultant drawings exhibited greater overall density variation and statistically significant local areas of density variation not seen in drawings where the niche construction mechanism was turned off. We (subjectively)

judged these niche constructed images to be more artistically appealing than the non niche constructed images.

For the experiments described in this paper, we extended the models of the previous section, employing a two-species engineering web. Previous studies had show this system is capable of forming ‘quasi-stable spatial structures with distinctive topologies’ [25]. The environment consists of a two-dimensional drawing surface populated by mobile drawing agents who are free to move over the surface of the canvas. Agents modify the canvas by drawing on it. The drawing forms the environmental materials that agents, through their actions, can modify. Additionally, a global depletion process causes the image to fade away over time if no agents work to maintain or change it.

Two different species of agent populate the canvas. Species A favours a niche with high local image entropy, species B favours low local image entropy. This preference is genetically encoded as a monoallelic locus, A , along with several other loci that determine the agent’s behaviour on the canvas. The genetic structure for both species is otherwise identical. The fitness, f , of both species depends on the entropy of the image over a small region surrounding the agent according to a simple linear function: $f = 1 - H_p + A(2H_p - 1)$, where H_p is the image entropy in the area around the point p , the agent’s current position. In the current experiments A may take the value 1 (species A) or 0 (species B).

A set of loci control agent motion and behaviour, which includes preferred speed and rate of turn, intensity and width of the lines drawn, and attraction to or repulsion from neighbouring agents of the same or opposing species. A normalised real number represents the allele at each locus.

A fixed number of randomly initialised agents populate the canvas and proceed to draw on it according to their genetically derived drawing and grouping behaviour. After an initial period, necessary to establish drawing patterns on a blank canvas, tournament selection is applied to randomly selected individuals of the same species. The fitter individual is allowed to breed, its offspring probabilistically replacing the weaker individual proportionate to the difference in fitness. Cross-breeding between species is not permitted, however a low background mutation rate grants the possibility of mutation of the A locus, meaning an agent can potentially mutate from one species to another. As with our studies of Section V, two different methods of tournament selection were tested: global and local. In each case offspring are placed in their parent’s region, the idea being that parents construct a heritable niche suitable for their offspring.

We ran a number of tests, comparing results between the niche constructing version and a neutral evolution version, where the probability of any agent replacing another is fixed at 0.5 (see Figure 8). In the niche construction version, agents of the same species group together to form quasi-stable regions of high or low image entropy, and indeed the system consistently generates images with regions of pronounced entropy differences not seen in neutral evolutionary runs. Cooperative behaviour between agents is necessary to successfully maintain

such differences. The longevity and stability of niches is more pronounced in the runs with spatial selection, following the principle that parents construct heritable environments for their offspring.

Due to the interdependence within related populations of agents in achieving specific-entropy environments, sub-populations with the same preference but different strategies also emerged. Lone invaders from one such sub-population would behave inappropriately in the environment produced by a different species, disrupting rather than reinforcing the high entropy. While this may reduce the fitness of the existing agents, the outcome of the interaction is typically worst for the fitness of the invader. This consideration gives rise to local regions where different behavioural strategies achieved a similar local entropy level. Analysis of the allele distribution confirms the visual and behavioural differences between different species. For example, a common strategy is for high entropy preferring agents to collectively make small, fast, interleaved circular marks, whereas low-entropy agents make large, slow and broad marks all of a similar colour.

Also of interest is the change in behaviour based on agent density. With few agents – a low average density – separate, stable niches are readily constructed in separate areas and there is little interaction between species. Populations sometimes fixate to low entropy preferring, as there are insufficient interactions for high entropy agents to form a stable niche. Movement of populations can result in interference between different entropy preferring populations, disrupting the system momentarily until new stable regions are formed.

The results of this study support the principles of boundary stability discussed in the previous studies, but not the principle of boundary emergence. The latter was harder to establish in this case, possibly because the interdependence of individuals in sub-populations, and the greater range of influence of each individual (because they are mobile), limits the freedom of breakaway groups to drift.

VII. CONCLUSION

In this paper we have set out a general approach to specifying how ecosystemic processes can be applied to creative domains, and have illustrated, using simple systems, how natural boundaries between constructed environmental conditions can emerge through drift, and can be reinforced through the inability of invading agents to bring about environmental change quickly enough before being out-competed. An essential requirement for this to occur is a spatial model with local rather than global competition. Additionally, individual agents must be able to modify their environment, and have preferences for particular environmental conditions. Boundary reinforcement is also shown to be fragile under certain conditions, such as in the 2D experiment where curved boundaries become eroded through the greater niche-constructing ability of the population on the outside of the curve. This simple set of studies sets out to illustrate some of the challenges of achieving emergent heterogeneity in an easily understood and readily deployable form, in which agents evolve preferences for specific niches

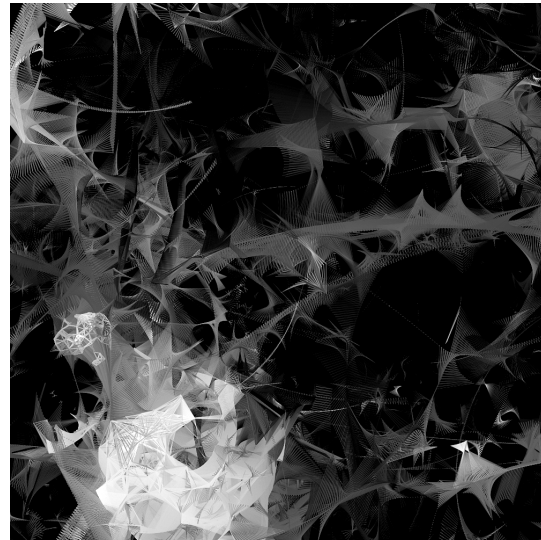
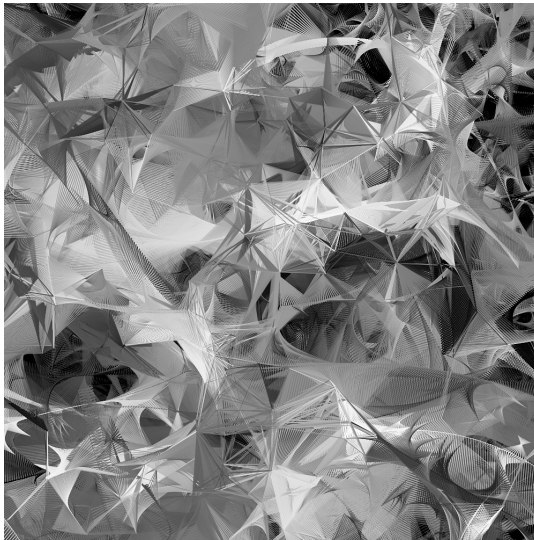


Fig. 8. Evolutionary dynamic line drawing sketch. Left: a sample image with random genetic drift (offspring have a fixed probability of 0.5 of replacing their parents). Right: the same system with niche construction. Offspring replace their parents probabilistically based on the suitability of their parents niche, tied to matching their entropy preference and the image entropy of their local environment. Drawings using this niche construction technique exhibit greater local variation and the emergence and maintenance of patterns and behaviours not seen in the neutral evolutionary version.

formed of perceptual features. It illustrates the viability of working with creative domains in the manner outlined in Section IV. The introduction of further methods established in the ecosystem modelling literature should be easily managed without undermining this paradigm.

ACKNOWLEDGMENTS

This research was supported by an Australian Research Council Discovery Project Grants, DP0877320 and DP1094064. We also thank the anonymous reviewers.

REFERENCES

- [1] F. J. Odling-Smee, K. N. Laland, and M. W. Feldman, *Niche construction: The neglected process in evolution*. Princeton, USA: Princeton University Press, 2003, no. 37.
- [2] C. Jones, J. Lawton, and M. Shachak, "Organisms as ecosystem engineers," *Oikos*, vol. 69, no. 3, pp. 373–386, 1994.
- [3] M. Conrad and H. H. Pattee, "Evolution experiments with an artificial ecosystem," *Journal of Theoretical Biology*, vol. 28, no. 3, pp. 393–409, 1970.
- [4] J. Brewster and M. Conrad, "Computer experiments on the development of niche specialization in an artificial ecosystem," in *Proceedings of the 9th annual conference on Genetic and evolutionary computation (GECCO 2007)*. ACM, 2007, pp. 301–307.
- [5] D. L. DeAngelis and W. M. Mooij, "Individual-based modeling of ecological and evolutionary processes," *Annu. Rev. Ecol. Evol. Syst.*, vol. 36, pp. 147–168, 2005.
- [6] E. A. Di Paolo, "Social coordination and spatial organization: Steps towards the evolution of communication," in *Fourth European Conference on Artificial Life*, P. Husbands and I. Harvey, Eds. Cambridge, MA, USA: MIT Press, 1997, pp. 464–473.
- [7] M. Silver and E. Di Paolo, "Spatial effects favour the evolution of niche construction," *Theoretical Population Biology*, vol. 70, no. 4, pp. 387–400, 2006.
- [8] M. Rönkkö, "An artificial ecosystem: Emergent dynamics and lifelike properties," *Artificial Life*, vol. 13, no. 2, pp. 159–187, 2007.
- [9] J. McCormack, "Eden: An evolutionary sonic ecosystem," *Advances in Artificial Life, Proceedings of the Sixth European Conference, ECAL*, vol. LNCS 2159, pp. 133–142, 2001.
- [10] J. McCormack, "Artificial ecosystems for creative discovery," in *Proceedings of the 9th annual conference on Genetic and evolutionary computation (GECCO 2007)*. ACM, 2007, pp. 301–307.
- [11] R. Suzuki and T. Arita, "How spatial locality can affect the evolution of niche construction," in *Proceedings of Artificial Life X*, 2006, pp. 452–458.
- [12] R. Axelrod, *The Evolution of Cooperation*. Basic Books, 1984.
- [13] M. Whitelaw, *Metacreation: art and artificial life*. Cambridge, Mass.: MIT Press, 2004.
- [14] C. W. Reynolds, "Flocks, herds and schools: A distributed behavioural model," in *Computer Graphics (SIGGRAPH '87 Conference Proceedings)*, vol. 21, no. 4, 1987, pp. 25–34.
- [15] T. Blackwell and M. Young, "Self-organised music," *Organised Sound*, vol. 9, no. 2, pp. 137–150, 2004.
- [16] E. Driessens and M. Verstappen, "Natural processes and artificial procedures," in *Design by Evolution: Advances in Evolutionary Design*, ser. Natural Computing Series, P. F. Hingston, L. C. Barone, and M. Zbigniew, Eds. Springer Verlag, 2008, ch. 5, pp. 101–120.
- [17] A. Eldridge and A. Dorin, "Filterscape: Energy recycling in a creative ecosystem," in *Proceedings of EvoWorkshops 2009*, M. Giacobini et al., Eds., vol. LNCS 5484, 2009.
- [18] J. McCormack, "Enhancing creativity with niche construction," in *Artificial Life XII*. Cambridge, MA: MIT Press, 2010, pp. 525–532.
- [19] G. R. Greenfield, "Co-evolutionary methods in evolutionary art," in *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, J. Romero and P. Machado, Eds. Springer-Verlag: Heidelberg, Germany, 2008, pp. 357–381.
- [20] A. Dorin, "A survey of virtual ecosystems in generative electronic art," in *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, P. Machado and J. Romero, Eds. Springer-Verlag: Heidelberg, Germany, 2008, pp. 289–309.
- [21] M. Hemberg, U.-M. O'Reilly, A. Menges, K. Jonas, M. d. C. Gonçalves, and S. R. Fuchs, "Genr8: Architects' experience with an emergent design tool," in *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, J. Romero and P. Machado, Eds. Springer-Verlag: Heidelberg, Germany, 2008, pp. 167–188.
- [22] A. J. Keane and S. M. Brown, "The design of a satellite boom with enhanced vibration performance using genetic algorithm techniques," in *Conference on Adaptive Computing in Engineering Design and Control 96*, I. C. Parmee, Ed. P.E.D.C., 1996, pp. 107–113.
- [23] J. McCormack and O. Bown, "Life's what you make: Niche construction and evolutionary art," in *Applications of Evolutionary Computing: EvoWorkshops 2009*, 2009.
- [24] R. L. Day, K. N. Laland, and J. Odling-Smee, "Rethinking adaptation: the niche-construction perspective," *Perspectives in Biology and Medicine*, vol. 46, no. 1, pp. 80–95, 2003.
- [25] M. Silver and E. D. Paolo, "Spatial effects favour the evolution of niche construction," *Theoretical Population Biology*, vol. 70, pp. 387–400, 2006.