

Niche Construction and the Evolution of Complexity

Tim Taylor

Institute of Perception, Action and Behaviour, University of Edinburgh
JCMB, The King's Buildings, Mayfield Road, Edinburgh EH9 3JZ, U.K.
tim.taylor@ed.ac.uk

Abstract

An individual-based model of the process of niche construction is presented, whereby organisms disturb the environment experienced by their neighbours. This disturbance in local conditions creates a niche that potentially could be filled by another species (which would then create still more niches and so on). The model is unique in allowing the complexity of the organisms—measured by the number of genes they possess in order to be well adapted to their local environment—to evolve over time, and is therefore the first model with which it is possible to study the contribution of niche construction to the evolution of organism complexity. Results of experiments demonstrate that the process of niche construction does indeed introduce an active drive for organisms with more genes. This is the first explicit example of a model which possesses an intrinsic drive for the evolution of complexity.

Introduction

For more than half a century, engineers, computer scientists and biologists have tried to recreate the dynamics of biological evolution in synthetic systems, ranging from computer simulations to evolving RNA molecules *in vitro*. It is a surprising fact that none of these attempts has succeeded in producing ongoing evolutionary dynamics in which innovations continue to arise, or where the complexity of the system, given any suitable measure, can be said to increase in an unbounded fashion.

The artificial life platform Avida has recently been used to study the evolution of complex features (Lenski et al., 2003), but the observed increase in the complexity of organisms arose only because a series of nine progressively more complex reward functions was added to the environment. After the organisms had evolved solutions to these functions, the system reached a fairly stable end point.

Studies of evolution *in vitro*, such as Spiegelman's and Orgel's experiments with evolving RNA sequences using a viral enzyme (Orgel, 1979), have also demonstrated the need for a better theoretical understanding of these issues. Reflecting on Orgel's results, Maynard Smith comments: "In this simple and well-defined system, natural selection does not lead to continuing change, still less to anything that

could be recognized as an increase in complexity: it leads to a stable and rather simple end point. This raises the following simple, and I think unanswered, question: *What features must be present in a system if it is to lead to indefinitely continuing evolutionary change?*" (Maynard Smith, 1988, p.221, emphasis added)

The theoretical biology literature provides some suitable starting points for answering Maynard Smith's question. Waddington recognised the need to develop a framework that described the logical structure of open-ended evolution, and published a suggestion of what such a framework might look like (Waddington, 1969). He went as far as to call this characterisation a new paradigm under which biological evolution should be studied.

The overall scenario is summarised as follows: "The complete paradigm must therefore include the following items: A genetic system whose items (Q s) are not mere information, but are algorithms or programs which produce phenotypes (Q^* s). There must be a mechanism for producing an indefinite variety of new Q^* s, some of which must act in a radical way which can be described as 'rewriting the program.' There must also be an indefinite number of environments, and this is assured by the fact that the evolving phenotypes are components of environments for their own or other species. Further, some at least of the species in the evolving biosystem must have means of dispersal, passive or active, which will bring them into contact with the new environments (under these circumstances, other species may have the new environments brought to them). These environments will not only exert selective pressure on the phenotypes, but will also act as items in programs, modifying the epigenetic processes with which the Q s become worked out into [Q^* s]." (Waddington, 1969, p.120)¹

This demonstrates that Waddington was fundamentally interested in situations where organisms constitute part of the environment experienced by other organisms. This leads to *niche construction*, whereby the presence of one species may introduce new niches in which other species can flourish.

¹The final symbol in the original text is printed as Q^* s, but this appears to be a typographical error.

ish (which in turn may introduce yet more niches and so on). Waddington predicted that such a system would produce a continually increasing number of species, and continually increasing phenotype complexity (*ibid.* p.119).²

Little work has been devoted to exploring this proposal, presumably because of the difficulties in capturing it fully with an analytical model. One of the motivations behind the work described here was to build a individual-based model that complied with Waddington's paradigm. This could then be used to test his predictions about the evolution of complexity and diversity.

In addition to the work with Avida already mentioned, a small number of other relevant models have recently appeared in the literature. Walker described a model specifically based upon Waddington's ideas of niche selection (but not niche *construction*) (Walker, 1999). However, in this case the potential niches were predefined and were *not* dependent on the presence of particular species (i.e. inter-organism interactions were not modelled). This was a serious flaw, and the results showed that it generally did not lead to the evolution of complex organisms.

Odling-Smee and colleagues have studied the evolutionary consequences of niche construction using more traditional population genetic models (Laland et al., 1999; Odling-Smee et al., 2003). They have shown that niche construction is a potent evolutionary force and can lead to unusual dynamics. In the current context, though, a disadvantage of such models is that they are only capable of describing the change in frequencies of a fixed number of genes over time. They are unable to model the introduction of new genes during the course of evolution, and are therefore not ideal for studying questions relating to possible increases in community diversity and organism complexity.

A different approach has been taken by Christensen and colleagues, who present a model in which individuals are described as vectors in a genotype space, and interactions between species are described by a pre-generated interaction matrix (Christensen et al., 2002). Their model produced dynamics in which relatively stable communities of species (which they refer to as quasievolutionary stable strategies, or q-ESS) exist for extended periods, separated by short transition periods of hectic reorganization. Furthermore, the average duration of the q-ESS periods increased slowly over time, in agreement to analysis of the fossil record. Despite the many attractive features of this model, it does not include a concept of organism complexity, and therefore cannot be used to address questions regarding how this evolves.

Finally, Pachepsky and colleagues described a spatially-explicit, individual-based model in which mutualistic relationships could evolve between organisms (Pachepsky et al.,

2002). They found that the possibility of such relationships lead to increased community diversity and stability. In agreement to the results of (Christensen et al., 2002), the composition of the communities tended to show periods of stability separated by short transition periods leading to new relatively stable states. However, the model was again not designed to address the evolution of organism complexity. It was also limited by the fact that mutualistic relationships could only arise between pairs of organisms rather than more extended ecological webs.

In the following sections, a new model is described which has been specifically designed to study the possible consequences of niche construction for the evolution of community diversity and organism complexity. The model overcomes many of the shortcomings of previous work in addressing these questions.

Description of the Model³

Before presenting a formal description of the model, it is worth giving a general overview of its design. It is an individual-based model in which the local environment experienced by an organism is expressed as an arbitrary vector of real numbers. An organism's behaviour, and the degree to which it is of adaptive value in the local environment, is abstractly modelled by requiring the organism to define a mathematical function which should closely match this "environment vector".⁴ Reproductive success is proportional to the closeness of match over a defined subset of the vector. The crucial part of the model is that the presence of an organism causes a perturbation in the environment vector experienced by all organisms in the local neighbourhood, thereby producing an environmental niche that potentially could be filled by a different species. As organisms can evolve arbitrarily complicated functions in order to match their environment, the model can therefore be used to study the role of niche construction in the evolution of organism complexity.

The formal definition of the model is as follows. Space is represented as a discrete two dimensional grid with wrapped boundaries in both dimensions (so the topology is toroidal). Each grid position, or *patch*, is denoted P_{ij} , where i and j are its spatial coordinates. A patch can contain zero or more *organisms*, up to a number limited by the local density parameter T_M . Each patch provides a local environment, which is experienced by all organisms located within it. This is represented as a vector $\mathbf{E}_{ij} = (E_{ij}^1, E_{ij}^2, \dots, E_{ij}^{L^E})$, where E_{ij}^k is a real number in the range $[0, 1]$. These elements E_{ij}^k can be interpreted as observables or attributes of the physical environment (e.g. temperature, humidity, levels of a nutrient, etc.). The local environment \mathbf{E}_{ij} can be influenced by the

³The source code of the model is available at <http://homepages.inf.ed.ac.uk/timt/papers/ncec/>

⁴This method of representing an environment, and the way in which organisms generate functions to match the environment, were inspired by and adapted from (Rocha, 2001).

²The sentence about "rewriting the program" shows that Waddington was also concerned with the problem of how fundamentally novel phenotypic traits can arise during evolution, which is another topic of major importance to artificial life research.

presence of organisms in a local neighbourhood of patches P_{xy} , where $x \in \{i - R_N, i + R_N\}$, $y \in \{j - R_N, j + R_N\}$, and R_N is the neighbourhood radius (a parameter of the model). The details of an organism's influence on the environment are described later. In the absence of any organisms in the local neighbourhood, $\mathbf{E}_{ij} = \mathbf{E}^A$, which represents the virgin abiotic environment.

An organism O_n is defined by the triplet (S_n, \mathbf{V}_n, D_n) which can be interpreted as, respectively, its sensitivity to particular environmental attributes, its preferred operating values for those attributes, and the manner in which it alters the environment through its behaviour. A schematic diagram of the structure of an organism is shown in Figure 1, and of a local collection of patches in Figure 2.

S_n is a pair (B_n, L_n) with $1 \leq (B_n, L_n) \leq L^E$. This specifies that the organism is sensitive to all environmental attributes in the range $[E^{B_n}, E^{B_n+L_n})$ (indices on \mathbf{E} wrap around such that $E^{L^E+1} = E^1$ and so on). This information is used in the calculation of an organism's fitness in a given environment, to be described later.

\mathbf{V}_n (the second element of O_n) is a vector of *genes*, $\mathbf{V}_n = (G_n^1, G_n^2, \dots, G_n^{C_n})$, of variable length $|\mathbf{V}_n| = C_n \geq 1$. The general idea is that the elements of \mathbf{V}_n represent transformation operations that can be applied sequentially to a default "preferences vector" \mathbf{P}^D (of length L^E equal to that of an environment vector \mathbf{E}) in order to generate a new vector \mathbf{P}_n representing the organism's preferred operating values for each of the attributes of the environment. The larger an organism's vector of genes \mathbf{V}_n , the more operations are applied to \mathbf{P}^D , and hence the more complicated the form of \mathbf{P}_n can become. The definition of a gene, and the way in which a vector of genes \mathbf{V}_n is used to generate a vector \mathbf{P}_n which can then be compared to the local environment vector \mathbf{E}_{ij} , follows the method described in (Rocha, 2001). For reasons of space, only a general outline of the method will be described here; for a full description, the reader is referred to Rocha's paper. A gene G_x is defined by the quintuple (F_x, p, s, r, \odot_x) . F_x is an element of \mathcal{F} , a small set of fuzzy set shapes each defined over a generic interval $[0, L^E]$. Similarly, \odot_x is an element of \mathcal{O} , a small set of fuzzy set operations. The action of a gene G_x on a given preferences vector \mathbf{P}_y is to apply the fuzzy set shape F_x to \mathbf{P}_y using the fuzzy set operation \odot_x . The elements p, s and r of G_x provide further details of the action, such as specifying the subset of elements of \mathbf{P}_y on which F_x is to be applied. See (Rocha, 2001) for details. For a given set of genes, the final form of \mathbf{P}_n is sensitive to the order in which the genes are applied, so even if we consider only organisms with a given, small number of genes, a large variety of different \mathbf{P}_n is possible.

D_n (the final element of O_n) is a quadruple of the form (F_n, p, s, r) , the elements of which have the same meaning as the corresponding elements of a gene. D_n determines the way organism O_n disturbs the local environment vector \mathbf{E}_{ij} . Specifically, the environment vector at a given patch P_{ij}




Figure 1: Structure of an organism. S_n specifies the environment attributes to which the organism is sensitive, and \mathbf{P}_n specifies its preferred values for those attributes. \mathbf{C}_{D_n} specifies the way in which the organism affects the local environment experienced by itself and other organisms.

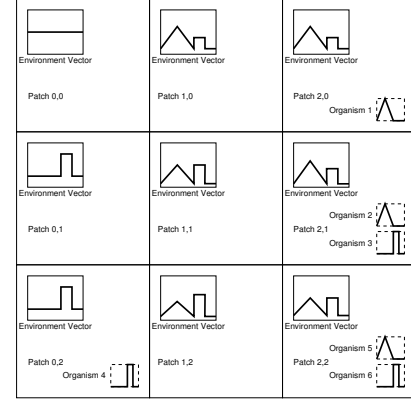


Figure 2: A local collection of patches. In this example the local neighbourhood radius R_N is 1, and all unshown patches neighbouring this local collection are assumed to contain no organisms. For each organism present, its environmental disturbance pattern \mathbf{C}_{D_n} is shown. Patch 0,0 displays the virgin abiotic environment because there are no organisms in its neighbourhood. The environment in all other patches has been modified by the presence of organisms.

at time t , $\mathbf{E}_{ij}^t = \frac{1}{|\mathcal{N}_{ij}|+1} (\mathbf{E}^A + \sum_{n \in \mathcal{N}_{ij}} \mathbf{C}_{D_n})$, where \mathcal{N}_{ij} is the set of organisms currently residing in all patches within the local neighbourhood radius R_N , and \mathbf{C}_{D_n} is the fuzzy set shape defined by D_n , given by F_n appropriately modified by p, s and r . In experiments where the model flag `ENV_DIST_RESTRICTED=true`, s is clamped at zero for all organisms, which limits the number of elements of \mathbf{C}_{D_n} which can take on nonzero values—that is, an organism can only disturb a limited subset of environmental attributes. This definition of \mathbf{E}_{ij}^t was adopted because it is additive with respect to the members of \mathcal{N}_{ij} , so a given collection of organisms will always produce the same disturbance regardless of the order in which their effects are considered.

The fitness $f_n(i, j)$ of an organism O_n in patch P_{ij} is defined as $f_n(i, j) = 1 - \frac{1}{L_n} \sum_{x=B_n}^{x=B_n+L_n-1} |\mathbf{E}_{ij}(x) - \mathbf{P}_n(x)|$ (where B_n and L_n are the components of S_n , as described earlier). The probability p_n^R of organism O_n reproducing at a given iteration of the simulation is given by $p_n^R = c^R f_n(i, j)$,

where c^R is a parameter of the model. Each organism also has a probability of death, p_n^D , associated with it. In some experiments this is kept constant, $p_n^D = c^D$, where c^D is a model parameter, and in others it is set to $p_n^D = c^D(1 - f_n(i, j))$.

Time in the model proceeds in discrete iterations. The top level algorithm is shown in Figure 3. The procedure `initialise` sets the environment vector of each patch to be \mathbf{E}^A , randomly generates a number of organisms (determined by the initial local density parameter T_0), each of size $|\mathbf{V}_n| = 1$, and places them in randomly chosen patches. `updateEnvironment` updates \mathbf{E}_{ij} for patch P_{ij} according to the organisms currently present within the local neighbourhood. `updateOrganisms` recalculates the fitness f_n and reproduction probability p_n^R (and, if appropriate, the probability of death p_n^D) of each organism within a patch, according to the current form of \mathbf{E}_{ij} . `killOrganismsStochastic` considers each organism O_n within a patch in turn, and kills it with probability p_n^D . Similarly, `reproduceOrganisms` makes a copy of each organism with probability p_n^R and places it in a randomly chosen patch within a dispersal radius R_D . When an organism reproduces, mutations may be introduced into the offspring: a randomly generated new gene is inserted into \mathbf{V}_n with probability p_+^M ; a randomly selected gene in \mathbf{V}_n is deleted with probability p_-^M (provided $|\mathbf{V}_n| > 1$); for each gene in \mathbf{V}_n , with probability $p_g^M/|\mathbf{V}_n|$, a randomly chosen element of the gene is mutated; similarly, with probability p_d^M , a randomly chosen element of D_n is mutated; and, with probability p_s^M , the value of the elements of S_n are independently altered by either -1 , 0 or $+1$. Finally, `checkOvercrowding` considers, for each patch, the mean number of organisms present in each patch within the local neighbourhood. If this exceeds the maximum local density threshold T_M , more organisms are killed from patches in the local neighbourhood (stochastically, according to their p_n^D values) until the density returns below T_M .

Experiments

Five sets of experiments were conducted in order to answer the question: *Does niche construction introduce an intrinsic drive for the evolution of organism complexity?* All experiments used the standard model configuration unless otherwise stated.⁵ Five experiments were conducted in each set, which had identical configurations apart from the seed supplied to the random number generator at run initialisation.

Set I was the standard trial set, using the default configuration unaltered. Set II was the main control set, in which `ENV_DIST_ENABLED=false` (i.e. the presence of organisms did not affect a patch's environment vector, so no niche

⁵This was: `ENV_SIZE=50x50`, `NUM_ITERATIONS=5000`, $R_N = 1$, $R_D = 2$, $T_0 = 0.1$, $T_M = 10$, $L^E = 100$. $\mathbf{E}^A = \mathbf{P}^D = (0.5, 0.5, \dots, 0.5)$, $p_+^M = p_-^M = p_g^M = p_d^M = p_s^M = 0.05$, $c^R = 0.5$, $c^D = 0.1$, $L_n = 20$, `DEATH_PROB_POLICY=constant`, `ENV_DIST_ENABLED=true`, `ENV_DIST_RESTRICTED=true`.

```

initialise

iteration = 0

while (iteration < NUM_ITERATIONS)
{
  if (ENV_DISTURBANCE_ENABLED)
    for each patch P; do
      updateEnvironment(P)
    end
  endif

  for each patch P; do
    updateOrganisms(P)
    killOrganismsStochastic(P)
    reproduceOrganisms(P)
  end

  for each patch P; do
    checkOvercrowding(P)
  end

  iteration = iteration + 1
}

```

Figure 3: The top level algorithm

construction could occur). Set III was similar to Set II, but the abiotic environment \mathbf{E}^A was more complex (see inset of Figure 4(f)) than the homogeneous version used in other experiments. This was designed to see to what extent any difference in dynamics between Sets I and II was due to the more complex environments present in Set I *per se*, rather than to the continuous process of niche construction. In Set IV, `ENV_DIST_RESTRICTED=false`, so organisms could affect an unlimited number of environmental attributes. Sustained evolutionary trends in this case were expected to be less likely, as ecologies were always open to the threat of a new species being introduced that had a devastating affect on the environment. Finally, in Set V, an organism's probability of death, $p_n^D = c^D(1 - f_n(i, j))$ [with $c^D = 0.5$], rather than being constant. Thus fitter organisms should tend to live longer as well as reproduce more frequently, effectively increasing the selection pressure for well-adapted organisms.

Results and Analysis

In all sets, the results obtained from each experiment of the set were qualitatively similar. In the following discussion, typical results from each set are highlighted.

The main results from Set I are shown in Figure 4(a)-(d). Figure 4(a) shows that the mean number of genes $|\mathbf{V}_n|$ of organisms in the population steadily grew over time, starting from 1 and reaching approximately 14–15 genes after 5000 iterations. The mean age of organisms in these runs (not shown) stayed fairly steady at about 3 iterations per organism, so 5000 iterations represents approximately 1667 generations. The overall population size (not shown) settled to a level of around 21500 individuals within the first 100 iterations, and stayed at that level for the rest of the run. The diversity of different species, where a species is

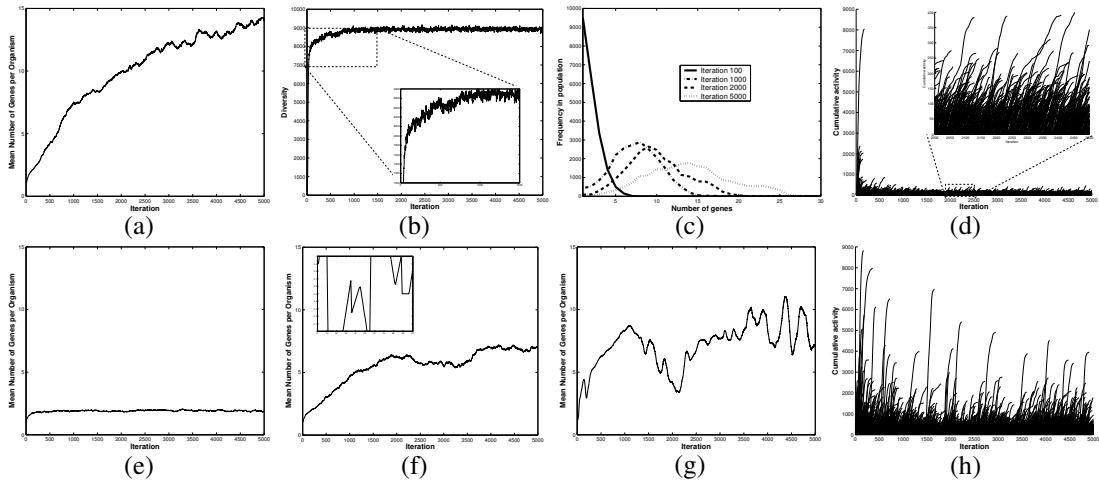


Figure 4: (a)–(d) Set I, Run 1: (a) Evolution of mean number of genes (b) Evolution of population diversity (c) Evolution of gene number frequencies (d) Species activity waves, (e) Set II, Run 1: Evolution of mean number of genes (f) Set III, Run 1: Evolution of mean number of genes (inset: \mathbf{E}^A used for Set III), (g)–(h) Set V, Run 4: (g) Evolution of mean number of genes (h) Species activity waves

defined as a collection of individuals that share exactly the same (S_n, \mathbf{V}_n, D_n) , also settled to a fairly constant value (of around 9000 species), although it only settled at this value after around 1000 iterations (Figure 4(b)). The inset of this figure shows that even after the population level had settled down after 100 iterations, the diversity of the system gradually increased for the next 900 iterations. However, the size of this increase was not large. Furthermore, if we look at the diversity of *significant* species (not shown), arbitrarily defined at a given iteration as those species that had 10 or more individuals present in the population, then the reverse trend is seen—the diversity of these species is around 300 at iteration 1, and gradually falls to around 200 by iteration 1000, after which time it remains fairly constant.

Figure 4(c) shows the distribution of genome sizes ($|\mathbf{V}_n|$) in the population at various points during the run. This demonstrates that the move to longer genomes is an *active* trend (McShea, 1994)—that is, it is not an artifact due to the presence of a lower limit of genome size, but indeed the whole population is moving towards higher genome sizes. Figure 4(d) shows an activity wave diagram of significant species throughout the run. Each line represents a single species, and is defined by the cumulative count of individuals of the species over time (Bedau and Brown, 1999). The inset of this figure demonstrates that many significant species coexist at any given time, and multiple species often tend to die out at around the same time (shown by the cessation of multiple activity waves)—although there is always a base level of active species that prevents the population from collapsing when such extinctions occur.

The change in genome size over time for a run from the control group in Set II is shown in Figure 4(e). In this case,

the mean size remains at around 1.8 genes for the duration of the run. In addition, the frequency distribution of genome sizes (not shown) does not change over time, but remains in a state similar to the plot for iteration 100 in Figure 4(c) for the duration of the run. These results demonstrate that there is no drive, either active or passive, for increased genome size in the situation where organisms do not cause disturbances in their local environment. This indicates that the active trend observed in Set I was indeed caused by niche construction.

In Set III, organisms again did not cause environmental disturbances, but the abiotic environment \mathbf{E}^A was more complex. In this situation, an increase in genome size was observed over time (Figure 4(f)), but this tended to tail off after reaching a level of around 7 genes. This shows that complex environments introduce selection pressure for complex organisms (i.e. with a large number of genes). However, in this case, the environmental complexity has been provided extrinsically—this is similar to the situation in (Lenski et al., 2003). In contrast, in Set I, we started with a simple environment, but the system displayed intrinsic dynamics for generating complex organisms (and environments).

The Set IV experiments involved a situation where organisms could again cause disturbances to the environment, but this time the breadth of the disturbance was unlimited. As expected, this tended to slow down the pace of evolution; the mean number of genes (not shown) did rise over the course of the runs, but more slowly than in Set I. In some runs it increased steadily, reaching a value of approximately 8 genes after 5000 iterations, whereas in other runs it seemed to peak at around 7 genes after around 3000 iterations and remain around that level of the rest of the run.

Perhaps the most surprising results, and the most variation between runs in a set, came from Set V, which had a similar configuration to Set I except that probability of death of an organism p_n^D was not constant, but rather inversely proportional to fitness. In these runs, the mean number of genes oscillated wildly throughout the run. An example is shown in Figure 4(g). The population diversity (not shown) was also somewhat lower than in Set I, in the region 6000–7000, and this also oscillated more than was observed in the other sets. However, the diversity of *significant* species was higher than in Set I, in the range 300–350 throughout the course of the run. Species activity waves were much higher than in other sets (Figure 4(h)). These results for diversity and activity waves are in line with the expectation that the new method of calculating probability of death would increase selection pressure—so the better adapted species tended to last for longer. The behaviour of the mean genome size still remains a little puzzling. One possibility is that because the population tends to comprise a higher proportion of highly adapted organisms that have been around for longer (which is confirmed by higher mean ages and mean probabilities of reproduction in this set), when one of these species finally does become extinct, a larger number of other species are likely to have depended upon it, and so the greater the impact on the stability of the ecosystem. Further experiments are underway to investigate this issue.

Discussion

Theoretical biologists have described situations in which there may exist drives for ongoing evolution rather than stasis, e.g. (Dawkins and Krebs, 1979; Van Valen, 1973), but the model presented here is the first to explicitly demonstrate a system with an intrinsic drive for the evolution of complexity. Organism-induced perturbations in the environment can be either beneficial or harmful to other organisms in the model. It therefore also provides a very general representation of symbiotic relationships.

It should be emphasized that, when an organism reproduces, the probability of mutation for each gene in \mathbf{V}_n is defined as $p_g^M/|\mathbf{V}_n|$. So the probability of a genome being copied with no mutations is roughly $(1 - p_g^M)$, regardless of length. If the mutation rate per gene was not normalized by $|\mathbf{V}_n|$, the model would soon reach the error threshold beyond which it is impossible to reliably transmit genomes to the next generation (Eigen and Schuster, 1977)—and this has indeed been verified experimentally. That is, the organisms are able to continually increase in complexity only because we are assuming they have found a way of reliably replicating long genomes, thereby avoiding the error threshold.

In terms of complying with Waddington's paradigm (see Introduction), the model still falls short. It successfully captures the concept of niche construction, but misses features such as the ability to fundamentally rewrite the genotype–phenotype mapping, and modelling epigenetic processes.

However, the results have demonstrated that niche construction by itself is sufficient to introduce an intrinsic drive for the evolution of organism complexity.

The question of how long this trend might continue in the model remains unanswered—in a couple of runs in Set I there were signs that the increase in genome length was starting to tail off towards the end of the run. Further work will investigate this, as well as factors that could prevent (or at least delay) any tailing off (such as allowing organisms to evolve different degrees of sensitivity to environmental attributes (i.e. allowing L_n in S_n to evolve), or even allowing the length of \mathbf{E}_{ij} to grow to simulate the evolution of fundamentally new types of behaviour). More work is also required to analyse the occupancy of different niches in the model, their stability, duration and evolution.

Finally, the environment described in the model is very abstract. A challenge for artificial life researchers is to understand how to build real (or artificial) environments which possess the same capacity for niche construction. The main challenge here is to build systems in which organisms constitute significant components of the environment experienced by other organisms. Some suggestions for progress in this direction are offered in (Taylor, 2004).

References

- Bedau, M. A. and Brown, C. T. (1999). Visualizing evolutionary activity of genotypes. *Artificial Life*, 5(1):17–35.
- Christensen, K., di Collobiano, S., Hall, M., and Jensen, H. (2002). Tangled nature: A model of evolutionary ecology. *Journal of Theoretical Biology*, 216:73–84.
- Dawkins, R. and Krebs, J. (1979). Arms races between and within species. *Proceedings of the Royal Society London*, B 205:489–511.
- Eigen, M. and Schuster, P. (1977). The hypercycle: A principle of natural self-organization. *Die Naturwissenschaften*, 64(11):541–565.
- Laland, K., Odling-Smee, F., and Feldman, M. (1999). Evolutionary consequences of niche construction and their implications for ecology. *Proceedings of the National Academy of Sciences USA*, 96:10242–10247.
- Lenski, R. E., Ofria, C., Pennock, R. T., and Adami, C. (2003). The evolutionary origin of complex features. *Nature*, 423:139–144.
- Maynard Smith, J. (1988). Evolutionary progress and levels of selection. In Nitecki, M., editor, *Evolutionary Progress*, pages 219–230. Univ. Chicago Press.
- McShea, D. W. (1994). Mechanisms of large-scale evolutionary trends. *Evolution*, 48(6):1747–1763.
- Odling-Smee, F., Laland, K., and Feldman, M. (2003). *Niche Construction: The Neglected Process in Evolution*. Princeton University Press.
- Orgel, L. (1979). Selection in vitro. *Proceedings of the Royal Society, London*, B, 205:435–442.
- Pachepsky, E., Taylor, T., and Jones, S. (2002). Mutualism promotes diversity and stability in a simple artificial ecosystem. *Artificial Life*, 8(1):5–24.
- Rocha, L. (2001). Evolution with material symbol systems. *Biosystems*, 60:95–121.
- Taylor, T. (2004). Redrawing the boundary between organism and environment. In *Proceedings of Artificial Life IX*. MIT Press.
- Van Valen, L. (1973). A new evolutionary law. *Evolutionary Theory*, 1:1–30.
- Waddington, C. (1969). Paradigm for an evolutionary process. In Waddington, C., editor, *Towards a Theoretical Biology*, volume 2, pages 106–128. Edinburgh University Press.
- Walker, R. (1999). Niche selection and the evolution of complex behavior in a changing environment—a simulation. *Artificial Life*, 5(3):271–289.