

“Niche Selection” and the Evolution of Complex Behavior in a Changing Environment—A Simulation

Richard Walker

Università di Palermo
Viale delle Scienze
Palermo, Italy
rwalker3@csi.com

Keywords

evolution, niche selection, neural networks, learning

Abstract One of the key problems in theoretical biology is the identification of the mechanisms underlying the evolution of complexity. This paper suggests that some difficulties in current models could be avoided by taking account of “niche selection” as proposed by Waddington [21] and subsequent authors [2]. Computer simulations, in which an evolving population of artificial organisms “selects” the niche(s) that maximize their fitness, are compared with a Control Model in which “Niche Selection” is absent. In the simulations the Niche Selection Model consistently produced a greater number of “fit” organisms than the Control Model; although the Niche Selection Model tended, in general, to produce organisms occupying simple niches, it was nonetheless more effective than the Control Model in producing well-adapted organisms inhabiting complex niches. It is shown that the production of these organisms is critically dependent on the rate of environmental change: Slow change leads to fit but undifferentiated populations, dominated by organisms occupying simple niches; differentiated populations, including well-adapted organisms living in complex niches, require rates of environmental change lying just beyond a mathematically well-defined critical value. In simulation “Niche Selection,” unlike conventional “Natural Selection,” provides a permanent selective bias in favor of simplicity. This tendency is counterbalanced by statistical forces favoring shifts from rare “simple niches” to commoner niches of greater complexity. Fit organisms inhabiting complex niches only emerge in conditions where the rate of environmental change is high enough to avoid the concentration of the population in very simple niches, but slow enough to permit step-by-step adaptation to niches of gradually increasing complexity. This result appears to be robust to changes in simulation parameters and assumptions, and leads to interesting conjectures about the real world behavior of biological organisms (and other complex adaptive systems). It is suggested that some of these conjectures might be relatively easy to test.

1 Introduction

One of the key problems in theoretical biology is the identification and analysis of the mechanisms underlying the evolution of complex organisms and behaviors. Molecular and paleontological work provides compelling evidence that complex organisms have

evolved step by step, with structural and behavioral adaptations to one set of environmental conditions being used as building blocks (preadaptations or exaptations) in subsequent responses to environmental change [9]. In the light of this work it seems intuitively obvious that a changing environment is a pre-condition for the emergence of complexity. It has proved remarkably hard, however, to incorporate this insight into formal models of biological evolution.

Biological adaptation is commonly modeled as “hill-climbing” on a static “fitness landscape” [22]. When environmental change is introduced into the picture [14], the fitness landscape is said to “deform.” This “deformation”—it is tacitly assumed—will affect the entire population. This point of view, while perfectly valid for small, geographically concentrated populations evolving over relatively short periods of time, makes this kind of model unsuitable for the study of long-term evolutionary trends, such as those we assume to underlie the evolution of complexity. In biological reality many successful species are geographically dispersed and inhabit a range of different environments. Waddington [21] and subsequent authors [2] have argued that just as the environment selects for “fit organisms,” populations of organisms may play a complementary and perhaps equally important role by “selecting” environmental conditions (“niches”) suited to their capabilities, reproducing faster in these niches than in others with less favorable characteristics. The biological reality of this phenomenon has been demonstrated experimentally in *Drosophila* [21]. In terms of Wright’s original concept, different niches (with their differing requirements) can be represented as different “regions” on an “extended fitness landscape.”

Environmental change implies continual changes in the requirements that the environment imposes on organisms. It can be plausibly argued that if a population of organisms can “choose” from a “population” of different “niches,” the probability that a subpopulation will find a habitat where it is viable, prolonging the evolutionary lineage, will be much higher than if every organism faces an identical, continually shifting, set of requirements. If this is true, “Niche Selection” is a vital component in the evolutionary process, playing a complementary role to “Natural Selection” by the environment.

This article presents the results of computer simulations designed to investigate this hypothesis. In the simulations, evolving populations of artificial organisms, subject to “Natural Selection,” were allowed to “select” the “niche(s)” best fitted to their capabilities. The results of the simulations not only confirm the evolutionary contribution of “Niche Selection,” but suggest that “Niche Selection” has important implications for the current debate on the evolution of complexity. As will be seen, “Niche Selection” tends, as a general rule, to favor simplicity over complexity. This effect, however, depends on the dynamic context. In slowly changing environments, “Niche Selection” produced populations entirely dominated by organisms with very simple behaviors; well-adapted organisms inhabiting complex niches emerged (as a small minority of the population) only when the rate of environmental change lay just beyond a closely defined “critical value.” The final sections of this article are devoted to the analysis of these results and to a discussion of their significance.

2 Static and Dynamic Fitness Landscapes—“Evolvable Fitness Formulas”

Research in the fields of genetic algorithms [11] and artificial life [15] has produced a broad range of models in which evolving populations of artificial “organisms” learn to perform a particular “task” via a process of random variation and “Natural Selection.” This work has shown that simulated “Natural Selection” often can find simpler, better performing solutions to problems than those suggested by an “engineering” approach [4]. The critical problems of classical artificial intelligence [6] nonetheless remain intact. Problem domains (e.g., a particular navigation task) are pre-defined by the ex-

perimeter and are often highly artificial; the architectures of the “organisms” used in experiments are tightly coupled to the specific application; scaling up to larger, more realistic domains remains problematic—in short current work has failed to produce the kind of complexity we observe in the natural world.

A number of researchers have attempted to address these problems. S. Kauffman and S. Johnson [13] have experimented with “coupled fitness landscapes” in which different “species” co-evolve. In these models the fitness landscape for an individual species is defined, not by the experimenter, but by the evolution of the other “species” (e.g., “prey” or “predators”) in the environment. Organisms are thus faced with a fitness landscape that changes over time. As in the paleontologists’ account of biological evolution, adaptations to the landscape at a particular time can act as pre-adaptations for the landscapes encountered at later stages of evolution.

Another approach is that of Lund and Parisi who have introduced the concept of “evolvable fitness formulas” [17]. In standard work with genetic algorithms, the formula used to define the fitness of individual organisms is defined by the experimenter. Lund and Parisi, on the other hand, treat the fitness formula as a variable, heritable attribute of the organism, subject to the same selective pressures as other attributes. In their model, populations of organisms “select” the “fitness formula” that gives them the highest fitness scores: they do not limit themselves to optimizing their performance on a particular task; they “choose” the task.

The way this model is presented in Lund and Parisi’s article is sometimes confusing. In particular it is never really clear how “fitness formulas” differ from ordinary phenotypic traits, for example, a larger stomach allowing an organism to benefit from efficient foraging behavior. It is relatively easy, however, to interpret the model in such a way as to avoid this difficulty.

In biological reality, “fitness formulas” are associated not with organisms but with “niches,” that is, sets of structural and/or behavioral requirements. Organisms tend to live in the same niches as their parents. It follows that, in a loose sense, “niche choices” are heritable. As Waddington has shown, preferences may be genetically determined [21]. In many cases, however, this form of “inheritance” will depend on non-genetic mechanisms. These include geographical inertia (the organism is unable to travel over large distances), imprinting (the organism imprints on some attribute of the parents’ environment to which it is exposed shortly after birth), and cultural transmission (parents transmit their own environmental preferences to offspring). In general, organisms change niches not because they have mutated but because of exogenous environmental change (e.g., changes in climate, invasion of the habitat by new organisms, changes in population density).

This interpretation of evolvable fitness formulas maintains the key characteristics of Lund and Parisi’s model: Fitness formulas are (indirectly) associated with organisms (via niche choice); they can be inherited; they are subject to variation. If all this is true there is no reason why they should not be subject to Darwinian mechanisms of selection. In this light Lund and Parisi’s computational expedient of seeing them as part of the phenotype is perfectly reasonable. When a population “selects” a fitness formula, what it is doing in practice is selecting the niche it will inhabit. In short Lund and Parisi have created a formal model for Waddington’s process of “Niche Selection.” This model provided key inspiration for the work reported in this paper.

3 The Model

In the simulations reported later in this article, individual organisms were modeled as Artificial Neural Networks (ANNs) with a fixed number of “input neurons,” a single output neuron, a variable number of “hidden neurons” (zero in the initial pop-

ulation), and arbitrary neuron to neuron “connections” (including feed-back as well as feed-forward connections). Connection strengths were set to random values, uniformly distributed between -0.1 and 0.1 . The architecture of the networks was coded explicitly. There was thus no attempt to make a principled genotype/phenotype distinction.

The output of the network was defined as the state of the output neurons after the state of each neuron in the network had been computed for a predefined number of iterations (three in the simulations reported here).

The requirements of different “niches” were modeled as “computational tasks” that organisms could compute to a measurable degree of accuracy. Each organism in the initial population was assigned an individual, heritable “task” that indirectly defined its “fitness formula” (see below). For the purposes of the model, organisms with different tasks were said to inhabit different “niches.” Tasks were modeled as randomly generated Boolean functions coded as bit strings where b_i defines the desired state of the output neuron for the four-dimensional input vector represented in the subscript (as written in binary notation). Individual tasks thus were coded in terms of the desired state of the output neuron for each of the possible states of the input neuron. With four input neurons there were $2^{16} = 65,536$ possible tasks.

As in much current work [16], organisms in the population were allowed to “adapt” not only genetically, over evolutionary time, but also via “learning” during their own lifetime. This, it was hoped, would enable them to adapt to changes in the environment without architectural modification. Given that current algorithms for training hidden neurons have no known biological counterpart [5] and are extremely expensive in computational terms [12], it was decided to adopt a more realistic and economic approach, reminiscent of early work on Perceptrons [18], [19]. Hidden neurons were treated as a repository of “feature detectors,” which were not subject to training. The output neuron was trained using the “generalized delta rule” [20].

At the end of training the population was sorted into three equal-sized “competence” classes. Competence was given by:

$$competence = \frac{N_{correct}}{N_{trials}}$$

where $N_{correct}$ is the number of correct computations achieved in a given number of trials and N_{trials} is the number of trials. Random performance by the model produces values of approximately 0.5.

Organisms’ fitness (i.e., their reproductive success) was defined to be a function of competence. For simplicity’s sake the model assumed that reproduction was asexual. Organisms in the class with the highest competence levels were made to produce two “clones”; those in the second class produced one; organisms in the third class were eliminated, ensuring that the total size of the population remained constant.

During cloning, offspring inherited their parent’s niche and neural architecture (including the strengths of connections prior to training). Lamarckian inheritance was not allowed. “Mutations” were introduced via the addition and/or destruction of neurons (with probability 0.2 of adding/destroying a single neuron in a single organism in one generation) and connections (with probability 0.1 of adding/destroying a single connection to a single neuron in one generation). The strength of new connections was set at random using the same distribution used for the original connections. The possibility of adding neurons and connections meant it was theoretically possible to evolve organisms of arbitrary size, architecture, and complexity.

Environmental change was simulated by randomly flipping bits in the string representing the learning task; differing rates of change were represented by varying the

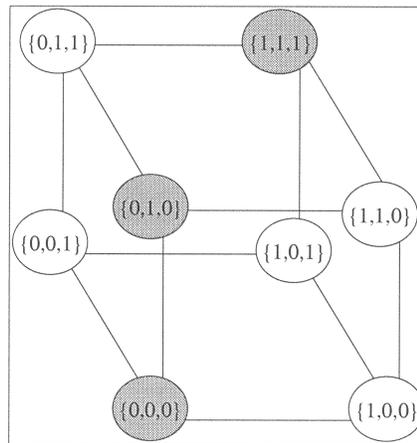


Figure 1. Representation of a three-dimensional Boolean function.

probability, P_{bf} , at which individual bits were flipped. The probability P_H that exactly H bits will be flipped is given, for $H > 0$, by:

$$P_H = (1 - P_{bf})^{2^D - H}$$

where D is the dimension of the task.

It follows that for the four-dimensional tasks used in simulation the probability, P_{stable} , that an organism will maintain the same task as its parent (i.e., $H = 0$) is given by:

$$P_{stable} = (1 - P_{bf})^{16}$$

In all simulations the cycle of training/selection/cloning and random variation was repeated for 100 generations, using a population of 99 individuals. In the initial population all organisms were assigned to a single niche of complexity four. Each organism in the population was trained for 600 learning cycles.

4 Measurement of Task Complexity

The use of abstract Boolean functions rather than the more naturalistic learning tasks used in much Alife research made it possible to quantify the complexity of individual tasks.

The “algorithmic complexity” of a task is defined by the length of the shortest possible coding for an algorithm capable of resolving the task [3]. For the purposes of this work a rough measure of the algorithmic complexity of individual functions was constructed by considering each bit in the bit string representing the function as lying at the vertex of a hypercube with coordinates $\{\dots 0, 0\}$ for the first bit, $\{\dots 0, 1\}$ for the second bit, and so on.

In Figure 1 we see a representation of the function $\{1, 0, 0, 0, 0, 1, 0, 1\}$, that is, of a function that maps the input $\{0, 0, 0\}$ to an output of 1, the input $\{0, 0, 1\}$ to an output of 0, and so forth.

Given this mode of representation, any Boolean function can be coded using a data structure including the value of a single b_k in the initial bit string, and a list of “edges”

$\{b_i b_j\}$ where $b_i \neq b_j$. The size of this structure sets an upper bound for the complexity of the function. Given that the only variable-length element in the coding is the list of edges, the complexity of the function is proportional to the number of edges in the list. Using order notation [8] we therefore can write:

$$C = \mathbf{O} \left(\sum_{i=1}^{2^D} \sum_{j=1}^{2^D} \text{Count}(b_i \neq b_j) \right), b_j \in \text{connected}(b_i)$$

where C is complexity, D is the length of the input vector (in bits), and $\text{connected}(b_i)$ is the set of vertices directly connected to b_i . In the case of the function illustrated in Figure 1, this gives $C = 7$.

Simulation results (data not shown) demonstrate a strong negative correlation between competence and task complexity ($r^2 = 0.98$). For complexity values in the range 0 to 16 the relationship between the two variables is very close to linear. These results appear to confirm the usefulness of the measure as an indication of the actual difficulty of performing specific computational tasks within an ANN architecture.

5 Statistical Distribution of Task Complexity

The tasks, used for simulation, map every possible combination of values for the input vector onto the value of the output neuron. The complexity measure used counts the number of times that neighboring input vectors require different outputs. The probability mass function is thus the binomial distribution:

$$\Pr(\text{Complexity} = m) = C_m^n p^m (1 - p)^{(n-m)}$$

where n is the number of “edges” on the hypercube, that is, $n = 2^{(D-1)}D$, and D is the dimension of the input vector.

For the randomly generated functions used in simulation $p = 0.5$. For the family of Boolean functions in four variables, maximum complexity is 32, minimum complexity is 0, and mean complexity is 16. For large n this distribution tends to the Gaussian.

It is obvious that the tasks used in simulation differ in essential ways from tasks in the real world. In particular physical constraints ensure that real-world input vectors show a high degree of correlation between input values to neighboring sensors; considerations of continuity ensure, what is more, that similar inputs require (in the majority of cases) similar outputs. This kind of constraint is lacking in the model. It is possible, however, to simulate its presence by restricting allowable input vectors to physically realistic values and by reducing the value of the p parameter in the distribution (i.e., by increasing the probability that neighboring input vectors will produce similar outputs). Although these changes lead to a reduction in mean task complexity they do not affect the form of the complexity distribution which, for large n , remains Gaussian. In brief, there are good theoretical reasons to expect “real world” distributions of task complexity to have the same form, if not the same parameters, as the distribution used for simulation.

6 Simulation Methodology

To test the hypothesis that “Niche Selection” facilitates the evolution of well-adapted organisms inhabiting niches requiring complex behavior, the model described above was compared with an otherwise identical Control Model in which all organisms were assigned the same task. This single task was subjected to random modification in the

Table 1. Competence and task complexity in the Niche Selection Model.

Complexity	Competence				
	0.00–0.49	0.50–0.79	0.80–0.89	0.90–0.99	1.00
0–4	0.0%	8.6%	5.0%	1.1%	17.0%
5–8	1.2%	9.7%	0.5%	0.9%	0.0%
9–12	11.0%	11.9%	1.6%	1.7%	0.0%
>12	25.5%	4.2%	0.0%	0.0%	0.0%

same way, and with the same probability, as the multiple tasks used in the Niche Selection Model. In the Niche Selection Model, the possibility of selecting among multiple niches allowed populations to improve their fitness both via conventional “Natural Selection” and via “Niche Selection.” Vice versa the single niche of the Control Model forced populations to rely for adaptation exclusively on Darwinian “Natural Selection.”

In preliminary testing simulations were repeated for many different values of the control parameters. The key tendencies observed were found to be robust with respect to relatively large changes in population size, number of generations observed, and number of learning cycles, as well as to modifications in the learning parameters, to changes in the probabilities used for mutations affecting the neural architecture, and to the complexity of the niches assigned to organisms at the beginning of each simulation.

During these investigations it was discovered that the overall behavior of the Niche Selection Model depended on the rate of environmental change, as modeled by the P_{bf} parameter. In order to take account of this dependency the Niche Selection and Control Models were compared in 60 paired simulation runs. In each simulation pair, niches in the two models were subjected to the same rate of “environmental change.” The first pair of simulations was conducted with $P_{bf} = 0.0025$. In each subsequent simulation the rate of evolutionary change was increased by 0.0025, enabling the simulations to explore the parameter range.

7 Simulation Results

In general the Niche Selection Model produced more competent organisms, performing simpler tasks than those in the Control Model. Mean competency for the Niche Selection Model was 0.60 ($S.D. = 0.27$) as compared to 0.32 ($S.D. = 0.19$) for the Control Model; mean task complexity was only 9.07 ($S.D. = 5.82$) as against 15.72 ($S.D. = 3.25$). The high standard deviations for competency and complexity in both models implied that the differences between means were not immediately significant. Subsequent investigation showed, however, that the high standard deviations were due to the effect of evolutionary change over time (in early generations complexity and competency values were lower than in later generations). When the competency and complexity figures were disaggregated by generation, it was seen that in every single generation except the first the Niche Selection Model produced a higher mean level of competence and a lower mean level of complexity than the Control Model. The difference in performance between the two models thus can be considered significant at any reasonable level of confidence.

Tables 1 and 2 show the aggregate distribution of competency and organisms for all organisms produced by the models during the 60 simulation pairs.

As can be seen from Table 1 a large proportion of organisms (17%) in the Niche Selection Model occupied niches of $0 \leq \text{complexity} \leq 4$. No organism in the Control Model (see Table 2) attained this degree of simplicity. Vice versa a very high proportion

Table 2. Competence and task complexity in the Control Model.

Complexity	Competence				
	0.00–0.49	0.50–0.79	0.80–0.89	0.90–0.99	1.00
0–4	0.0%	0.9%	0.2%	0.1%	0.0%
5–8	0.4%	1.7%	0.2%	0.2%	0.0%
9–12	8.9%	8.5%	0.7%	0.6%	0.0%
>12	72.1%	5.5%	0.0%	0.0%	0.0%

Table 3. Number of competent organisms in subpopulations with tasks of homogeneous complexity.

Complexity	Competence ≥ 0.9		Competence ≥ 0.8	
	NS	CM	NS	CM
0–4	107,422	503	137,202	1,733
5–8	5,493	1,036	8,484	2,256
9–12	10,245	3,666	19,772	7,802
>12	0	0	4	48

NS = Niche Selection Model, CM = Control Model

of organisms in the Control Model (72.1%) had a competence of less than 0.5 and occupied niches with *complexity* ≥ 12 ; in the Niche Selection Model these organisms occupied a much smaller proportion of the population (25.5%).

A comparison between sub-populations with tasks of homogeneous complexity found that, although there was no statistically significant difference between the two models in the mean levels of competence achieved (data not shown), the Niche Selection Model consistently produced a higher number of competent organisms than the Control Model (see Table 3).

This effect applied not only to niches of low complexity, where the Niche Selection Model was obviously favored by its tendency to produce large numbers of very competent organisms, performing very simple tasks, but also to niches with *complexity* ≥ 8 . As is apparent from Table 3, this superior performance of the Niche Selection Model was robust with respect to changes in the definition of what constitutes a “competent organism.” The simulation data thus support the initial hypothesis that Niche Selection effectively facilitates the evolution of organisms well-adapted to niches requiring complex behavior.

The evolutionary dynamics produced by the Niche Selection Model also differed radically from those of the Control Model. In simulations with “Niche Selection,” the mean competence of the population climbed rapidly during early generations until it reached a plateau. Once the plateau was reached, competence remained relatively stable. Task complexity followed a similar pattern, with a short period of initial adjustment leading to a stable equilibrium after a few generations. In the Control Model, on the other hand, competence and task complexity oscillated violently from generation to generation. Faster rates of environmental change lead to increased rates of oscillation. Figures 2 and 3 compare the dynamics of the Niche Selection and the Control Models in a single simulation pair with $P_{bf} = 0.06$.

As noted previously the behavior of the Niche Selection Model was found to be critically dependent on the rate of environmental change. Figures 4–8 show variations in key variables for gradually increasing values of P_{bf} . All values refer to the complete set of organisms produced over the whole simulation. Inspection of these figures shows, for the Niche Selection Model, three separate regimes of behavior. In the Control Model the distinction between these regimes is blurred or absent. In what

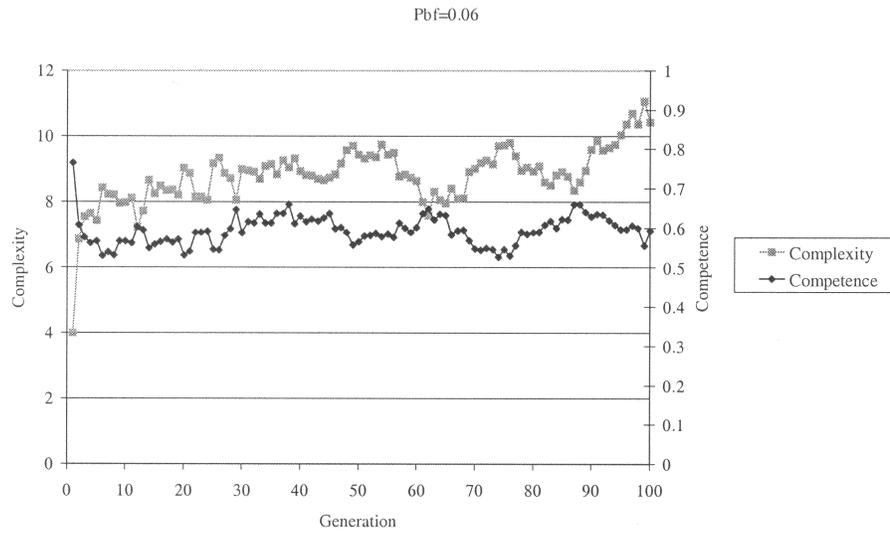


Figure 2. Evolution of competence and complexity in a single simulation with the Niche Selection Model ($P_{bf} = 0.06$).

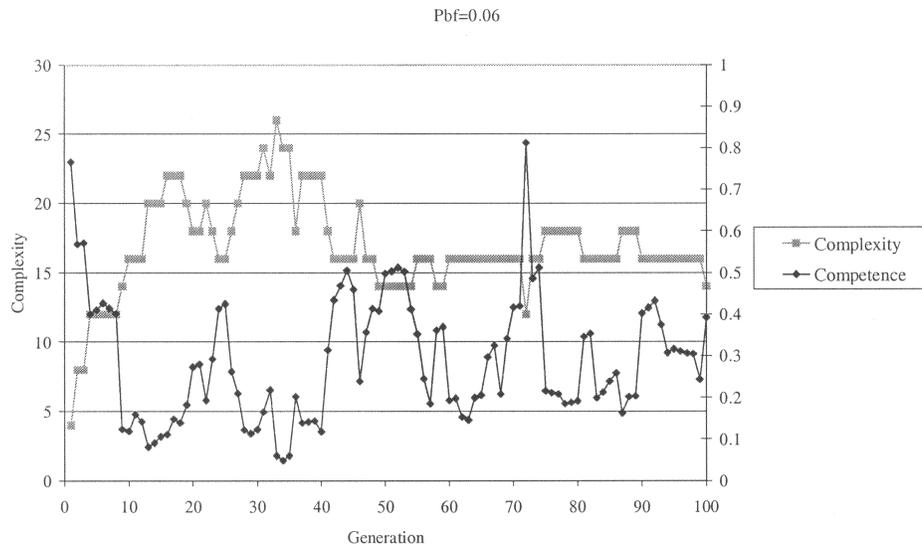


Figure 3. Evolution of competence and complexity in a single simulation with the Control Model ($P_{bf} = 0.06$).

follows there will be no attempt to identify distinct performance regimes for the Control Model.

7.1 Range I: Slow Environmental Change—Undifferentiated Populations of Very Competent, Simple Organisms

In simulations with $0.00 \leq P_{bf} \leq 0.04$, the mean competence achieved by the Niche Selection Model during simulation was high (0.96 for $P_{bf} = 0.0025$, 0.79 for $P_{bf} = 0.04$). The mean complexity of assigned tasks was low, however (0.66 for $P_{bf} = 0.0025$, 3.81 for $P_{bf} = 0.04$). Variance in complexity was initially low though rapidly

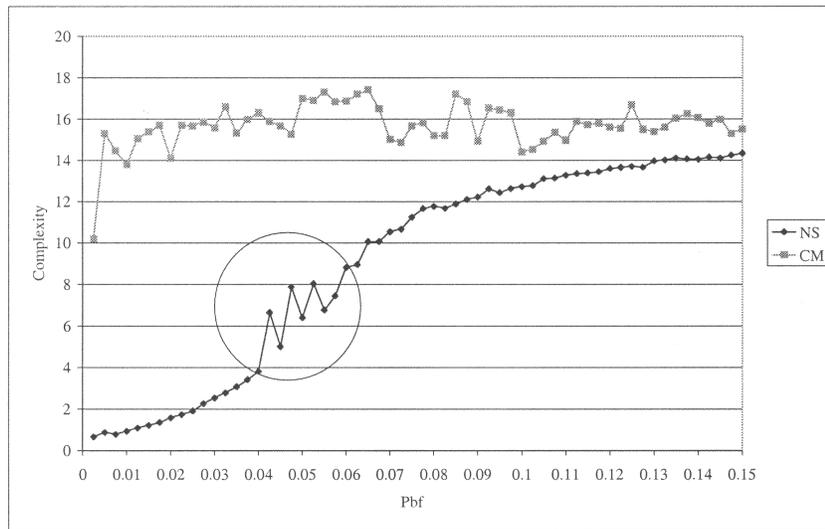


Figure 4. Variations in mean complexity for gradually increasing values of P_{bf} .

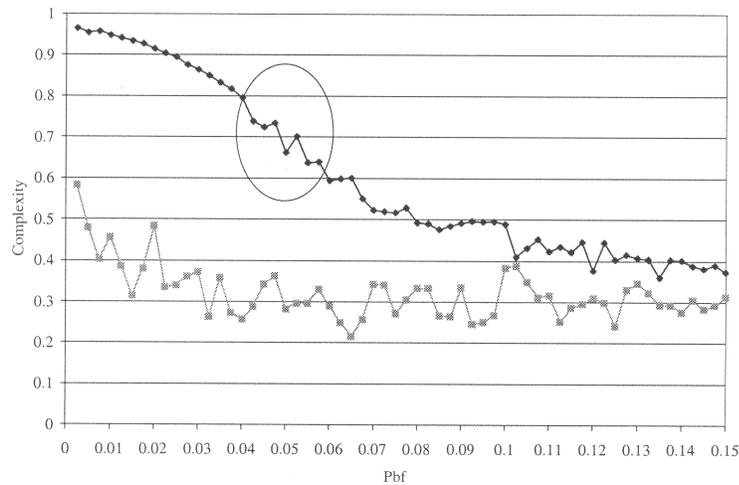


Figure 5. Variations in mean competence for gradually increasing values of P_{bf} .

rising (2.36 for $P_{bf} = 0.0025$, 10.76 for $P_{bf} = 0.04$) with a strong tendency for the population to concentrate in a single niche of zero complexity (97% of the population for $P_{bf} = 0.0025$, 30.1% for $P_{bf} = 0.04$). No simulation produced a single organism with *competence* ≥ 0.8 and *task complexity* ≥ 12 . Only very small numbers (0.05% of the population) were produced with *competence* ≥ 0.8 and *task complexity* ≥ 8 . Tests with a model in which “Niche Selection” was maintained, but learning inhibited (data not shown), suggested that in this range learning was making no significant difference to the complexity or the competence of the organisms produced by the model.

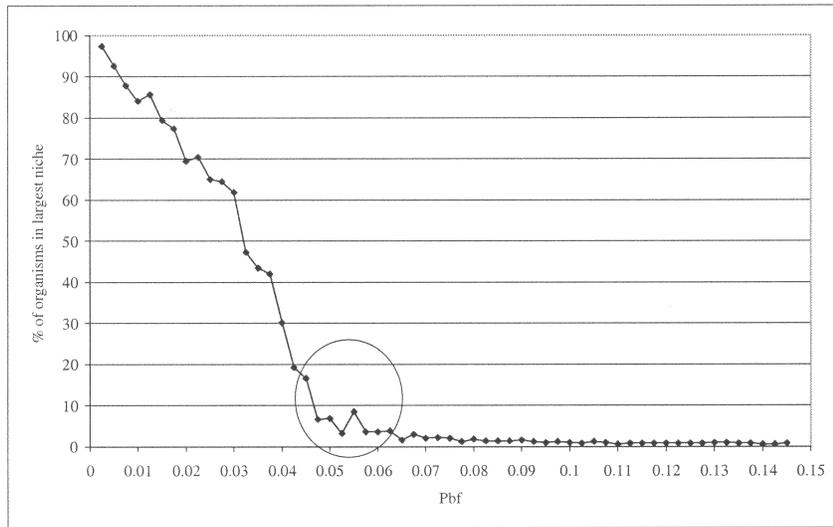


Figure 6. Percentage of organisms in largest niche—Niche Selection Model. In the Control Model all organisms belong to the same niche.

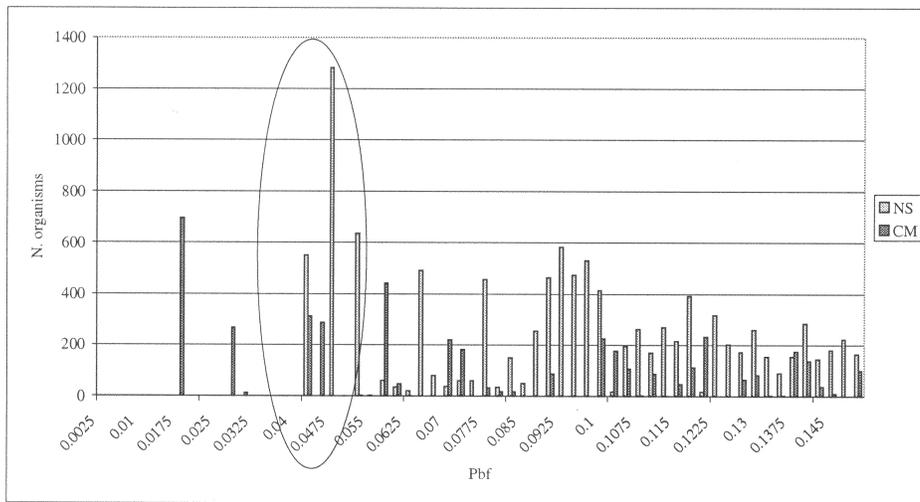


Figure 7. Variation in number of organisms with *competence* ≥ 0.9 and *complexity* ≥ 10 for gradually increasing values of P_{bf} .

7.2 Range 2: Intermediate Levels of Environmental Change—High Differentiation with Small Numbers of Competent Organisms Occupying Complex Niches

Between $P_{bf} = 0.04$ and $P_{bf} = 0.06$ (the circled area in Figures 4–8) the behavior of the Niche Selection Model changed radically. Although 7.4% of the population continued to be made up of organisms occupying niches of *task complexity* ≤ 4 , simulations in this range registered a dramatic increase both in the mean complexity of tasks (from 3.81 for $P_{bf} = 0.04$ to 8.80 for $P_{bf} = 0.06$), and in the variance of complexity (which for $P_{bf} = 0.0475$ reached 16.2). The population diversified rapidly,

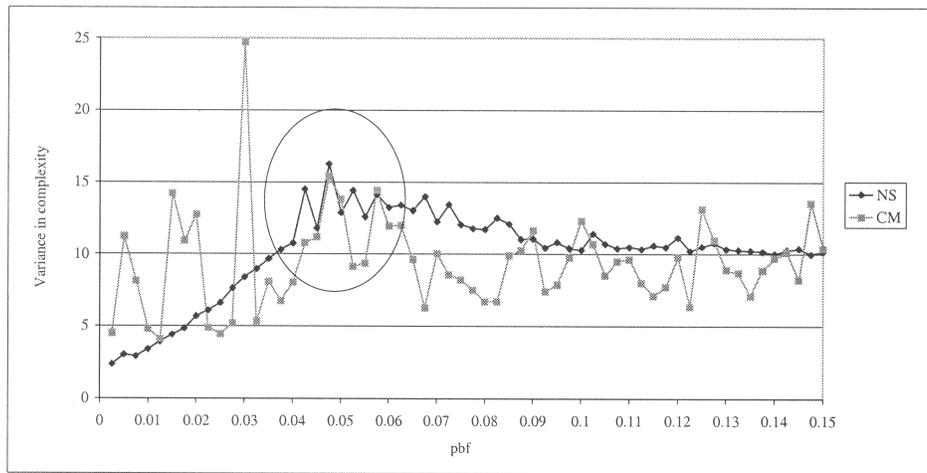


Figure 8. Variance in complexity for gradually increasing values of P_{bf} .

with the percentage in the most populated (low complexity) niche dropping from 30.1% for $P_{bf} = 0.04$ to 3.6% for $P_{bf} = 0.06$. Although mean competence fell from 0.79 to 0.59, the most competent organisms continued to achieve a competence score of 1.00. A small but significant proportion of the population succeeded in achieving high levels of competence even in relatively complex niches, with the highest rate (12.9% of the population with $competence \geq 0.9$ and $task\ complexity \geq 10$, 3.4% with $competence \geq 0.95$ and $task\ complexity \geq 12$) being registered in the simulation with $P_{bf} = 0.0475$. High “success rates” (more than 5% of the population with $competence \geq 0.9$ and $task\ complexity \geq 10$) were also registered for $P_{bf} = 0.0425$ and for $P_{bf} = 0.0525$. In other simulation runs, however, “success rates” were much lower. It was observed that for P_{bf} in this range, the number of competent organisms in complex niches depended critically on the choice of seed for the random number generator used by the simulator. This reflected an extreme sensitivity to initial conditions that was absent for values of P_{bf} in Range 1 and Range 3 (see Figure 7).

Comparison between the standard Niche Selection Model and an alternative model in which learning was inhibited (see Figure 9) showed that for P_{bf} in this range (circled in the figure), the former model produced significantly higher mean task complexity levels than the latter. There was no significant difference in the mean levels of competence produced by the two models. Investigation of individual lineages showed shifts both from high to low complexity tasks and vice versa. In most cases the observed shifts were small, involving a single bit flip. They were also relatively rare: Lineages sometimes remained in the same niche for up to 10 generations.

7.3 Range 3: High Levels of Environmental Change—Falling Numbers of Fit Organisms

Whereas simulation data showed a clearly marked transition between Range 1 and Range 2, the demarcation between Range 2 and Range 3, while real, was less well defined. It may be noted, however, that for $P_{bf} > 0.06$, simulations show a rapid reduction in the presence of organisms occupying low complexity niches. Though the mean complexity of tasks continues to increase (8.96 for $P_{bf} = 0.0625$, 12.71 for $P_{bf} = 0.10$, 14.34 for $P_{bf} = 0.15$), the rate of increase is much slower than in Range 2,

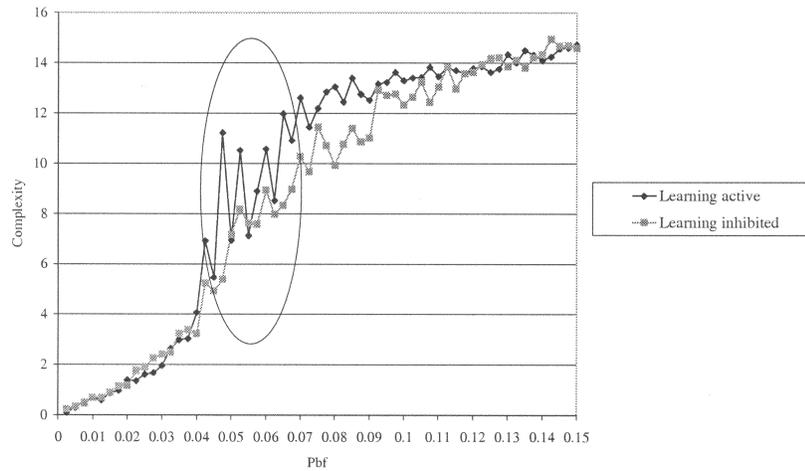


Figure 9. Mean complexity in the Niche Selection Model (last five generations), with learning activated or inhibited.

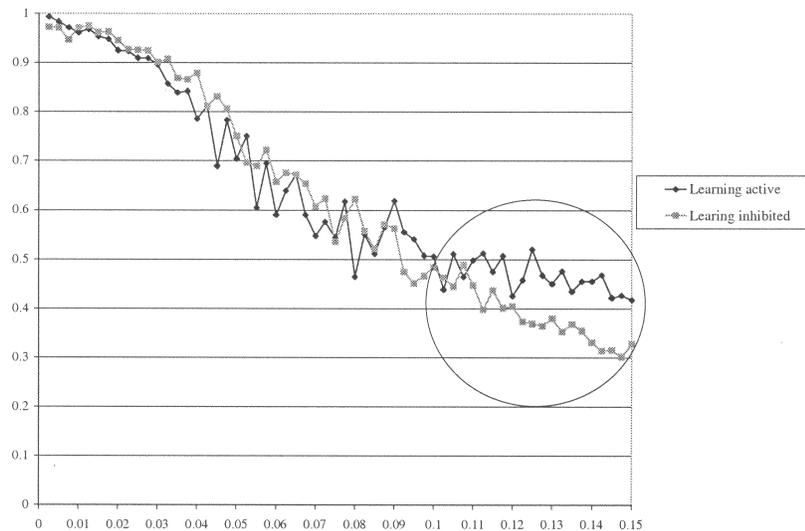


Figure 10. Mean competence in the Niche Selection Model (last five generations), with learning active or inhibited.

and variance in complexity (13.25 for $P_{bf} = 0.06$, 10.19 for $P_{bf} = 0.15$) is lower, falling with increasing values of P_{bf} . Mean competence declines in direct proportion to the increase in P_{bf} (0.59 for $P_{bf} = 0.06$, 0.50 for $P_{bf} = 0.10$, 0.42 for $P_{bf} = 0.15$). For $P_{bf} > 0.08$, even the best performing organisms fail to achieve a competence of 1.00. Simulations in this range continue to produce a number of competent organisms inhabiting complex niches. In no simulation, however, does the proportion of such organisms in the overall population achieve the levels characterizing Range 2. As values of P_{bf} increase, shifts from niche to niche become both more frequent and larger. The difference in complexity levels between simulations with learning and those in which learning is inhibited gradually disappears. In this regime learning appears to provide a fitness benefit that is absent for lower values of P_{bf} (see Figure 10).

8 Discussion

The main results of the simulations may be summarized as follows:

1. In sub-populations with homogeneous task complexity, the Niche Selection Model produces a greater number of competent organisms than the Control Model.
2. The population distributions produced by the Niche Selection Model are dominated by organisms inhabiting relatively simple niches (except for simulations with high rates of environmental change).
3. Despite 2., the Niche Selection Model produces a higher number of well adapted organisms inhabiting complex niches than the Control Model.
4. The qualitative behavior of the Niche Selection Model is critically dependent on the rate of environmental change. Low rates of environmental change lead to the domination of the population by very competent organisms occupying extremely simple niches. There exists, however, a critical rate of change above which the population diversifies rapidly. Rates of change just higher than this critical rate appear to be optimal for the evolution of well-adapted organisms occupying complex niches. As the rates of change increase beyond the critical level, there is a rapid fall in the proportion of competent organisms occupying complex niches.
5. In the critical regime of environmental change, learning appears to play a key role in encouraging the emergence of complex organisms. At rates of change below or a long way beyond the critical level, this role is absent.

It is useful to discuss these results one by one.

8.1 Higher Numbers of Competent Organisms Produced by the Niche Selection Model

It is proposed that the superior performance of the Niche Selection Model in the production of competent organisms is due to the greater range of options it provides to populations during their evolution. In conventional “Natural Selection,” the only way in which organisms can evolve is through modification of the phenotype. “Niche Selection,” on the other hand, provides a new dimension for evolution. This means that when a lineage reaches what, for “Natural Selection,” would be a “local optimum,” it can improve its fitness further by moving into a new niche. This is exactly the behavior recorded by the Niche Selection Model in simulations with low rates of environmental change. In these simulations the huge majority of organisms in the population rapidly concentrated in niches of extremely low complexity. In the Control Model, on the other hand, organisms were not allowed to “select” their own niches; the tasks they had to resolve were, on average, more complex than the tasks chosen by organisms in the Niche Selection Model (see Figure 4). It was inevitable that organisms allowed to choose low complexity niches would outperform organisms that did not have this freedom.

At higher rates of environmental change, the ability of populations in the Niche Selection Model to choose their own niche was used not so much to escape local optima as to guarantee survival.

At high rates of environmental change, organisms are faced continually with new challenges (in the model: changes in the task they have to perform). The loss in competence, suffered by a competent organism when its task changes, is roughly proportional to the Hamming distance between the bit strings representing the old and new tasks.

In this context, therefore, the ability of populations to select their own niche translates into the ability to select the niche that has changed least with respect to the niche where the population originally evolved. In this way some portion of the population is always able to maintain its fitness and pass on its accumulated adaptations to a new generation; “Niche Selection” thus exerts a stabilizing influence on niche occupation patterns. Simulation data (see Figure 2) confirms that even at relatively high rates of evolutionary change the Niche Selection Model manages to maintain relatively stable levels of complexity and competence.

This advantage of “Niche Selection” may be contrasted with the Control Model, where all organisms develop in the same niche. Here any kind of environmental change affects every organism in the population at the same time; populations are forced continually to restart the adaptive process from scratch; mean competence and task complexity levels fluctuate wildly (see Figure 3); adaptive gains are never accumulated. In simulation this explains the generally low levels of competence achieved; in the biological world it would imply a severe risk of extinction.

8.2 Domination of the Population by Organisms Occupying Simple Niches

The simulation data suggest that this result, which is interesting in its own right, may be relevant to the current debate concerning the evolution of complexity. “Niche Selection,” unlike “Natural” Darwinian selection, has a clear direction. While there will be endless disagreement concerning the selective advantages and disadvantages of increased phenotypic complexity, it is tautologically obvious that simpler niches are easier to live in than more complex ones. In brief, if similar organisms are randomly distributed among niches of differing complexities, organisms in niches with simple requirements will achieve, as a general rule, greater reproductive success than those inhabiting niches with more complex requirements. Data showing lower mean levels of task complexity in the Niche Selection Model with respect to the Control Model provides evidence that, in the right conditions, this tendency can lead to a systematic selective bias in favor of simplicity.

8.3 Relatively High Number of Well-Adapted Organisms Occupying Complex Niches

In the simulations, environmental change takes the form of changes in the task that organisms are asked to perform. In general, as just argued, “Niche Selection” will favor shifts from complex to simpler niches. It is also possible, however, that when an organism is pre-adapted to a particular niche (in particular when the organism already possesses a neural architecture capable of “learning” the task associated with the new niche), there will be little or no competence penalty for added task complexity. In these cases a random walk from niche to niche may actually lead to a net increase in task complexity. Even if each shift in favor of increased complexity is less probable than a shift in the opposite direction, it is still possible that a sequence of shifts, starting from a population occupying a relatively simple niche, will, at some stage, produce a population of well-adapted organisms occupying a relatively complex niche. In view of the relatively low probability of niche shifts in the direction of greater complexity, there is likely to be an upper limit on the size of the populations produced in this way. For high levels of task complexity this will always be small, a prediction supported by simulation data showing that very few competent organisms produced by the Niche Selection Model achieve a task complexity of 12 or more (a value lower than the mean value for the ensemble of all possible niches), and that the size of the populations in these relatively complex niches is extremely small (even if higher than those produced by the Control Model). This provides support for the idea, propounded by Gould [10], that complexity is an exceptional product of the evolutionary process rather than its inevitable outcome.

8.4 Critical Dependency on the Rate of Environmental Change

In the model simple niches are much rarer than more complex ones. Purely random environmental change (as opposed to selection-driven shifts from niche to niche) will lead, in general, to an increase in the complexity of the niches “offered” to a population. It can be shown that the critical dependency of the Niche Selection Model on the rate of environmental change is regulated by the balance between this “statistical” tendency and the countervailing selective bias introduced by “Niche Selection.” Consider the sub-population inhabiting the “simplest” niche in an environment. In order for this population to be stable, it is necessary that each organism in the niche produces an average of one offspring that continues to inhabit the same niche. As seen earlier, the probability that an offspring remains in the same niche as its parent is given by:

$$(1 - P_{bf})^{2^D}$$

where D is the “dimension” of the task. If organisms in the top competence class produce Q offspring, the equilibrium condition is given by:

$$Q(1 - P_{bf})^{2^D} = 1$$

This equation is closely related to Eigen and Schuster’s equation for genomic stability in the face of random mutation [7].

Solving the equation for the parameter values used in simulation ($Q = 2$, $D = 4$) gives $P_{bf} \cong 0.042$. Thus it may be predicted that for rates of environmental change below this “critical value” ($P_{bf} < 0.042$) each parent will, on average, produce more than one offspring occupying its own niche, causing the sub-population in the niche to increase until it dominates the overall population; this is exactly what was observed in Range 1 of the simulations, where a large proportion of the population is occupied by organisms in a single niche of complexity 0.

It can be predicted further that for P_{bf} values above the critical value each parent will, on average, produce less than one offspring in its own niche; in these circumstances no niche will be stably occupied for long periods. This prediction is supported by simulation data showing, for $P_{bf} \geq 0.0425$, a rapid increase in niche diversity and in the frequency with which organisms shift to niches different to those occupied by their parent. For rates of environmental change below the critical value, it has been shown that the population is almost entirely confined to niches of low complexity. In these conditions the emergence of fit organisms occupying complex niches is extremely improbable. For rates of change above the critical value, no population stably inhabits the same niche. It follows that in this regime organisms can “choose” between niches of differing degrees of complexity. They thus have the “option” of choosing complex niches. As has been seen already, the probability that an organism can adapt successfully to such a niche depends on its evolutionary history. The kind of history most favorable to the colonization of complex niches is one where the adaptations accumulated in one niche can act as pre-adaptations for subsequent stages in evolution, enabling a lineage to colonize a sequence of niches of gradually increasing complexity. Pre-adaptation is most likely when each niche in a sequence is as similar as possible to its predecessors. The probability that any particular shift from niche to niche will be “small” is a diminishing function of the rate of evolutionary change. It follows that higher rates of environmental change will correspond to lower levels of pre-adaptation and lower levels of competence for organisms attempting to occupy complex niches. If this is so, the optimal rate of environmental change for the production of well-adapted organisms inhabiting complex niches is a rate just above the critical value, that is the lowest

rate of change compatible with the presence of a pool of complex niches available to organisms. This prediction is fully confirmed by the simulation results.

8.5 Learning and the Emergence of Complexity

In the absence of learning a single neural architecture can compute only one task. As a result, for an organism capable of perfectly performing its assigned task, any change in task definition leads to a fall in competence exactly proportional to the Hamming distance between the old and the new task. With learning, on the other hand, a given neural architecture can adapt so as to compute any one of a broad set of different functions (in the model: the set of all linearly separable functions taking, as arguments, the output of the input and hidden neurons with direct connections to the output neuron). It follows that organisms that are able to learn may be able at times to switch from an old to a new task without any corresponding penalty in competence. Superior adaptability undoubtedly explains the higher mean competence levels achieved by Range 3 simulations in which organisms were allowed to learn compared to simulations in which this ability was inhibited (see Figure 10).

For lower values of P_{bf} (close to the critical level), learning does not appear to confer any competence advantage; simulations with learning do produce, however, a far higher number of competent organisms occupying complex niches than those without. This result can be explained by the observation that, for values of P_{bf} close to the critical value, learning populations can maintain their fitness by selecting tasks similar or identical to those to which they originally adapted (an option not open to populations living in conditions of very rapid change). In these circumstances learning does not improve the mean competence of the population (competence levels in simple niches are often very close to the theoretical maximum). What it does do is increase the range of options. Organisms, in other words, can maintain fitness either by staying in the niche to which they originally adapted or by moving to another, perhaps more complex, niche for which their neural architecture is “pre-adapted.” This explanation is supported by the high variance in task complexity (see Figure 8) in simulations with P_{bf} close to the critical value. It implies that without learning, the probability of producing competent organisms, capable of inhabiting complex niches, would be very low. Simulation data support this conclusion, showing that without learning, the model completely failed to produce organisms of *competence* > 0.8 and *task complexity* ≥ 10 and produced very few with *task complexity* ≥ 8 .

9 “Real World” Relevance

Like much Alife work the simulations presented in this paper include a relatively high number of free parameters. This implies that specific quantitative values produced by the model, in particular the “critical” value for the rate of environmental change, have no claim to validity outside the arbitrary parameter regime used for simulation.

This having been said, it should be recalled that the qualitative results of the simulations have been tested for sensitivity to the values chosen for specific parameters and have been found to be robust. It has been shown, what is more, that there are theoretical reasons to expect that the most important of the underlying assumptions used in constructing the simulations (the shape of the statistical distribution of task/niche complexity) is insensitive to the introduction of “real world” constraints. In these conditions it seems at least possible that the main qualitative results could be valid not just in simulation, but in the “real world” as well. The simulations can be used, in other words, as a source of qualitative conjectures that may be of interest to theoretical and observational biologists as well as to students of other kinds of complex adaptive systems. “Niche Selection,” the simulations suggest, creates a “selective bias” in favor of

simplicity. The way in which this bias operates depends on the rate of environmental change. In particular the model suggests the hypothesis that there exists a critical rate of change that encourages the emergence of well-adapted organisms occupying complex niches. This proposal is obviously reminiscent of recent work by many authors on “adaptation at the edge of chaos” [13] and “self-organized criticality” [1].

The main conjectures deriving from the model are, at least in principle, testable. Low rates of change should lead, the model predicts, to a gradual process of evolutionary simplification; rates of change just above a “critical value” should, vice versa, encourage complexity; extremely high rates of change should lead to mass extinction (as is trivially obvious).

The testing of these hypotheses against data on the evolution of biological organisms would require considerable ingenuity—in particular in the quantification of difficult concepts such as “the rate of environmental change” or “complexity.” It seems likely that testing might be easier in the case of other kinds of complex adaptive systems. The complexity of an ecosystem, for example, could be quantified in terms of species numbers; rates of environmental change could be measured by variance in climatic variables (measured on a suitable time scale) or by the frequency of “traumatic events” (such as fire). The complexity of national economies could be represented by the diversity of exports or of production; the rate of environmental change could be measured by variance in the prices for key inputs or exports. Given the existence of these (and similar) measures, it should be relatively easy to test the basic hypothesis that low rates of environmental change favor simplification while rates of change just above an observable critical value produce a rapid increase in complexity.

It is suggested that the preliminary evidence for the robustness of the results produced by simulation, the importance of the conjectures to which they give rise, and the practical possibility of testing these conjectures justifies deeper investigation of the issues raised by the interaction between “Niche Selection” and environmental change, both in future modeling and in discussion of its real-world implications.

References

1. Bak, P. (1996). *How Nature Works*. New York, NY: Springer-Verlag.
2. Belew, R. K., & Mitchell, M. (1996). Introduction. In R. K. Belew & M. Mitchell (Eds.), *Adaptive Individuals in Evolving Populations*. Reading, MA: Addison-Wesley.
3. Chaitin, G. J. (1975). Randomness and mathematical proof. *Scientific American*, 232, 47–52.
4. Clark, A. (1997). *Being There, Putting Brain, Body and World Together*. Cambridge, MA: MIT Press.
5. Crick, F. (1989). The recent excitement about neural networks. *Nature*, 337, 129–132.
6. Dreyfuss, H. L. (1972). *What Computers Can't Do, A Critique of Artificial Reason*. New York, NY: Harper and Row.
7. Eigen, M., & Schuster, P. (1978). The hypercycle part B: “The abstract hypercycle.” *Naturwissenschaften*, 65, 7–41.
8. Garey, M. R., & Johnson, D. S. (1979). *Computers and Intractability - a Guide to the Theory of NP-Completeness*. New York, NY: W. H. Freeman and Company.
9. Gould, S. J. (1991). Exaptation: A crucial tool for an evolutionary psychology. *J. Soc. Issues*, 3, 43–65.
10. Gould, S. J. (1996). *Full House*. New York, NY: Three Rivers Press.
11. Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: University of Michigan Press.

12. Judd, J. S. (1990). *Neural Network Design and the Complexity of Learning*. Cambridge, MA: MIT Press.
13. Kauffman, S. A., & Johnson, S. (1992). Co-Evolution to the edge of chaos: Coupled fitness landscapes, poised states and co-evolutionary avalanches. *Artificial Life II*, 325–369, Redwood City, CA: Addison-Wesley.
14. Kauffman, S. A. (1993). *The Origins of Order*. New York, NY: Oxford University Press.
15. Langton, R. A. (1989). (Ed.). *Artificial Life I*. Redwood City, CA: Addison-Wesley.
16. Littman, M. (1996). Simulations combining evolution and learning. In R. K. Belew & M. Mitchell (Eds.), *Adaptive Individuals in Evolving Populations* (pp. 465–477). Reading, MA: Addison-Wesley.
17. Lund, H. H., & Parisi, D. (1996). Pre-adaptations in populations of neural networks, evolving in a changing environment. *Artificial Life*, 2, 179–197.
18. Minsky, M., & Pappert, S. (1969). *Perceptrons*. Cambridge, MA: MIT Press.
19. Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65, 386–408.
20. Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart, J. L. McClelland, & PDP Research Group I (Eds.), *Parallel Distributed Processing* (pp. 318–364). Cambridge, MA: MIT Press.
21. Waddington, C. H. (1975). *Evolution of an Evolutionist*. Ithaca, NY: Cornell University Press.
22. Wright, S. (1931). Evolution in mendelian populations. *Genetics*, 16, 97.