

A Comparison of Speciation, Extinction, and Complexification in Neuroevolution With and Without Selection Pressure

Derek James
The University of Louisiana at
Lafayette
PO Drawer 43772
Lafayette, LA 70504-3772
(337) 482-1000
djames@gmail.com

ABSTRACT

In the study of biological evolution, neutral theories are invoked by some researchers to explain the dynamics of evolving populations regardless of selection pressure. The current study compares the dynamics of speciation, extinction, and complexification in two sets of populations of evolving artificial neural networks. One set of populations evolved under selection pressure, their survival dependent upon performance at a control task, while the other set of populations had survivors chosen randomly. Despite predictions to the contrary, the results showed significant differences in all three dynamics, suggesting that neutral models are incomplete explanations at best and that selection pressure constrains evolutionary search in very specific ways.

Categories & Subject Descriptors

I.6.0 SIMULATION AND MODELING, General

General Terms

Algorithms, Theory

Keywords

Speciation, extinction, complexification, neuroevolution, selection pressure, neutral theory

1. INTRODUCTION

The central question under investigation are these: To what extent (if any) does selection pressure influence the dynamics of an evolving population? Are the rates of speciation and extinction affected by selection pressure? Is the rate and range of complexity in an evolving population affected by selection pressure? In other words, if survivors for subsequent generations were chosen randomly, rather than based on some criteria, would we expect the evolving populations to exhibit the same characteristics as those under selection pressure?

Copyright is held by the author/owner(s).
GECCO '08, July 12–16, 2008, Atlanta, Georgia, USA.
ACM 978-1-60558-130-9/08/07.

2. METHODS

The algorithm used as the testbed for these ideas was NeuroEvolution of Augmenting Topologies (NEAT), a system for evolving the connection weights and topologies of artificial neural networks (ANNs). NEAT is ideal for investigating the effects of selection versus neutral conditions on such issues because it explicitly employs a method of speciation, subdividing the population in an attempt to diversify the evolutionary search. Just as in organic ecosystems, new species emerge and old species perish. Thus, NEAT provides a model of speciation and extinction dynamics at an appropriate level of detail.

In the selection pressure condition, each individual's fitness was a function of how well they performed on a traditional motor control task, double pole balancing (Figure 1).




Figure 1. Double pole balancing is a control task in which the system tries to keep two poles on a mobile cart from falling over by exerting force each step in either horizontal direction.

In this simulated version of the task, two poles of unequal length are attached (with a single degree of freedom) to a mobile cart which can move along a one-dimensional track in either direction. Each time step, the neural network receives information about the state of the system, such as the velocity of the cart and the angle of each pole. The network output indicates the directional force to apply to the cart. The goal is to keep both poles from falling and breaking the horizontal plane. This task initially used a single pole, with the addition of a second pole to increase difficulty. The task has been used as a reliable benchmark for control systems for many years.

3. RESULTS

3.1 Speciation

The mean speciation rate was calculated as the number of new species originating in each generation, averaged over all generations of all 10 runs in each condition.

Speciation rates were lower in selection pressure condition ($M = 6.49$, $SD = 1.33$) than in the random condition ($M = 7.63$, $SD = 0.98$). This difference was significant, $t(18) = 2.18$, $p < 0.05$. Figure 2 provides a graphical representation of this result.




Figure 2. Comparison of mean speciation rates between the selection and random conditions. Error bars represent standard error.

3.2 Extinction

The mean extinction rate was calculated as the number of species that had no surviving members from the previous generation.

Mean extinction rates were also lower in the selection condition ($M = 6.26$, $SD = 1.23$) than in the random condition ($M = 7.34$, $SD = 0.92$). This difference was also significant, $t(18) = 2.22$, $p < 0.05$. The results can be seen graphically in Figure 3 below.




Figure 3. Comparison of mean extinction rates between the selection and random conditions. Error bars represent standard error.

3.3 Complexity

The complexity of a network was simply taken as a measure of the number of genes in each chromosome.

Mean complexity for the final generation of the 10 runs was compared between conditions. Individuals in the pole balance condition ($M = 91.48$, $SD = 59.19$) were on average less complex than those in the random condition ($M = 180.18$, $SD = 109.44$). This difference was significant, $t(18) = 2.25$, $p < 0.05$. Figure 4 depicts this information graphically.




Figure 4. Comparison of mean complexity for all members of the final generation from the selection and random conditions. Error bars represent standard error.

Thus, there were significant differences in all the evolutionary dynamics observed between populations evolving with and without selection pressure, and in each case the mean values were lower in the selection than in the random condition.

4. CONCLUSION

What are the ramifications for the experimental results found here?

For the evolutionary biologists, the experiments described here suggest that neutral theories, at both the micro and macro level, insufficiently explain evolutionary dynamics of populations and inaccurately minimize the role of selection.

The researcher in evolutionary computation should strongly consider the use of a speciation mechanism in order to reduce bloat and efficiently constrain the search for candidate solutions. Even so, the work described here indicates that starting small and using niching mechanisms alone does not mitigate bloat. Researchers should consider performing sample runs without fitness as a baseline for comparison, and also using pruning or parsimony techniques either as part of the run or post hoc. Starting small and niching are useful techniques for mitigating superfluous structure, but they are not sufficient in and of themselves to eliminate bloat.

Further research remains to be conducted in understanding the complex interactions inherent in evolving populations. One conclusion is clear, that evolutionary dynamics such as those studied here are not subject to simple analysis or generalities, such as “fitness increases bloat” or “complexification always corresponds with increased fitness.” Such dynamics are subject to the particular features of the algorithm, and their behavior does not necessarily yield simple explanations.