

Adaptive populations of endogenously diversifying *Pushpop* organisms are reliably diverse

Lee Spector
Cognitive Science
Hampshire College
Amherst, MA 01002, USA
lspector@hampshire.edu

To appear in *Proceedings of the 8th International Conference on the Simulation and Synthesis of Living Systems, Artificial Life VIII*, December, 2002.

Conference URL: <http://parallel.acsu.unsw.edu.au/complex/alife8/>

ABSTRACT

This paper discusses the evolution of diversifying reproduction. We measured the average difference between mothers and their children, the number of species, and the degree of adaptation in evolving populations of endogenously diversifying digital organisms using the *Pushpop* system. The data show that the number of species in adaptive populations is higher than in non-adaptive populations, while the variance in the differences between mothers and their children is less for adaptive populations than for non-adaptive populations. In other words, in adaptive populations the species were more numerous and the diversification processes were more reliable.

Adaptive populations of endogenously diversifying *Pushpop* organisms are reliably diverse

Lee Spector

Hampshire College

Amherst, MA 01002

lspector@hampshire.edu

Abstract

This paper discusses the evolution of diversifying reproduction. We measured the average difference between mothers and their children, the number of species, and the degree of adaptation in evolving populations of endogenously diversifying digital organisms using the *Pushpop* system. The data show that the number of species in adaptive populations is higher than in non-adaptive populations, while the variance in the differences between mothers and their children is less for adaptive populations than for non-adaptive populations. In other words, in adaptive populations the species were more numerous and the diversification processes were more reliable.

Evolution of diversifying reproduction

Questions about the evolution of diversity have a long history, with many proposals having been made to account for the observed diversity of the biosphere (Gould 2002). For example, several authors have discussed how environmental variation and geographical separation can promote diversity, either under the influence of natural selection or without it (genetic drift). These discussions generally build on a foundational assumption that reproductive systems naturally and necessarily produce diversity upon which selection or drift can act. Questions about the evolution of diversity therefore depend on a more fundamental set of questions about the evolution of diversifying reproductive systems.

How do diversifying reproductive systems themselves evolve? Prior work on the origins of life has considered how quasi-reproductive processes in a primeval soup could become sufficiently *non*-diversifying to jump-start a truly reproducing lineage (Maynard Smith & Szathmary 1999). A common assumption for the aftermath of such an event is that natural selection will act to refine the reproduction process to improve replication accuracy within limits imposed by the reliability of the underlying chemistry. George C. Williams made the case as follows:

The fittest possible degree of stability is absolute stability. In other words, natural selection of mutation rates has only one possible direction, that of reducing the frequency of mutation to zero.

That mutations should continue to occur after billions of years of adverse selection requires no special explanation. It is merely a reflection of the unquestionable principle that natural selection can often produce mechanisms of extreme precision, but never of perfection. (Williams 1966, pp. 138–9).

Perfectly accurate reproduction would preclude evolution; Williams understood this and wrote that “evolution takes place, not so much because of natural selection, but to a large degree in spite of it” (p. 139). Williams’s view is in tension, however, with recent work on the evolution of genetic mechanisms. For example, the diversifying function of recombination is central to most theories of the evolution of sex, and Maynard Smith and Szathmary note that “every sexual population that has been carefully studied has been found to have genetic variability in the rate of recombination” (Maynard Smith & Szathmary 1999, p. 90). Recent work on microbiological mechanisms of diversification (e.g. gene duplication, shuffling, and horizontal gene transfer) reinforces the view that methods and degrees of diversification are themselves products of complex evolutionary histories and that an understanding of the origins and evolution of life will require an understanding of these histories.

Diversification of digital organisms

Populations of digital organisms (computer programs that run and reproduce in a simulated environment) can be used to study general features of evolutionary systems. For example, Ray’s *Tierra* system (Ray 1991) has been used to study the evolution of parasitism, and Adami’s *Avida* system has been used to study the “survival of the flattest” (that is, the survival of otherwise mediocre but reproductively robust replicators) in environments with high mutation rates (Wilke *et al.* 2001).

In most work on digital organisms diversification is exogenously controlled. In *Tierra* and *Avida* certain operations (for example code-copying) are automatically subject to random errors (mutations) that occur with a probability set by the experimenter. This provides a ready-made (and in fact unavoidable) diversification

mechanism: exact cloning is usually a simple procedure in these systems, and any organism that attempts to clone itself will automatically produce diverse offspring. It is possible for programs in these systems to diversify in other ways as well, but the external control of diversification limits the questions that one can ask about the evolution of the diversification mechanisms themselves. Most prior work also requires seeding with a hand-crafted “ancestor” replicator, further pre-determining the resulting reproduction and diversification mechanisms. An exception is Amoeba (Pargellis 2001), in which replicators are produced by spontaneous generation in a process that is nonetheless driven by exogenously controlled mutation.

The biological phenomena at issue in the present study can best be studied in systems within which replication and diversification strategies must emerge *ex nihilo*, and within which the degrees and types of diversification are endogenously controlled and thereby subject to selection and evolution. Such systems can be used to explore questions about how the nature of the emergent diversification mechanisms relates to other features of an evolving population, for example adaptation.

A large body of previous work has explored “self-adaptive” evolutionary computation in problem-solving contexts (see, e.g., (Stephens *et al.* 1998)). In most of these studies the algorithms for reproduction and diversification are fixed in advance although numerical parameters (such as mutation rates) are subject to variation and selection. Systems capable of more completely endogenous diversification include Edmonds’s tree-based “Meta-Genetic Programming” system (Edmonds 1998), Teller’s graph-based PADO system (Teller 1996), and the stack-based Pushpop system described below.

Methods

The experiments were conducted using the *Pushpop* system that evolves digital organisms expressed in the Push programming language (Spector 2001; Spector & Robinson 2002). Push is a high-level, stack-based language that supports flexible manipulation of multiple types of data including Push code. Pushpop is an *autoconstructive evolution system*, which is defined in (Spector & Robinson 2002) to be “any evolutionary computation system that adaptively constructs its own mechanisms of reproduction and diversification as it runs.” A more conventional genetic programming system that evolves Push programs (PushGP) has also been developed.

The Pushpop algorithm is best understood as a variant of a traditional genetic programming algorithm (Koza 1992). Evaluation of an organism’s fitness in Pushpop produces a numerical measure of the organism’s problem-solving ability and also a collection of potential children (Push programs). “Fitness” here, as in most genetic programming work, denotes the performance of

an individual on a computational problem that is used to determine survival; fitness causes reproductive success here, in contrast to conventional biological usage. Children are added to the following generation on the basis of fitness tournaments between their mothers — they themselves have not yet been tested for fitness, so their own fitness values cannot be used for this purpose. If there are insufficient children to populate the next generation (in which case the population is not yet “reproductively competent”) then randomly generated organisms are used. Most runs become permanently reproductively competent after at most a few hundred generations. We generally continue running Pushpop even after a solution has been found, in order to observe the changes in evolutionary dynamics that follow such an event. More information on Push, PushGP, and Pushpop, including source code for Push and PushGP, can be found at <http://hampshire.edu/ljspector/push.html>.

In this study we enforced a “no cloning” policy: children were never added to a population if they were identical (genotypically) to their mothers or to other children already in the population. Without such a policy, and without imposed mutations, perfect replicators quickly overwhelm the population, precluding future evolution. It is critical to note that the no cloning policy forces mothers to diversify their children but it does *not* constrain, to any significant degree, the manner or amount of diversification. In particular, minimal diversification strategies (e.g. those that change only a single, unused symbol from mother to child) are possible and do in fact emerge (as shown in the data below). The no cloning policy mandates that *some* diversification occur, but neither the methods nor the degree of diversification are pre-determined by the no cloning policy.

The Push language includes a rich library of code-manipulation instructions that can be used by organisms to construct their offspring. These instructions *allow* for the use of copy-error-based reproduction strategies of the sort used in Tierra and Avida, but their availability does not mandate the use of such strategies, pre-determine parameters for such strategies, or limit the ways in which such strategies can be combined with different reproductive methods.

Pushpop is capable of solving relatively difficult computational problems in comparison to other digital organism systems. Avida is usually assessed on simple 2-input Boolean problems, which is reasonable considering the low-level assembly language in which Avida programs are expressed. Tierra and Amoeba were not designed to be used in problem-solving contexts, so problem-based measures of the adaptation of a population are not available. With Pushpop we generally use standard test problems from the genetic programming literature, often involving numbers or other data types.

Parameters were chosen to produce roughly the same

number of adaptive and non-adaptive populations. The population size was 1024, the number of generations was 2048, the tournament size was 4, the fitness-conferring problem was integer symbolic regression (Koza 1992) of $y = 5x^2 + x - 2$ with 16 fitness cases ($x = 0$ to $x = 15$), the instruction execution limit was 64, the maximum number of points (symbols and pairs of parentheses) in a program was 64, and the full standard Push instruction set was available. In half of the runs we included instructions (NEIGHBOR, ELDER, and OTHER) that support sexual recombination of evolving organisms by allowing an executing organism to access the code of other organisms; see (Spector & Robinson 2002) for details.

We measured the average difference between mothers and their children, the number of species, and the degree of adaptation of each population each generation. The difference between two organisms was calculated as the sum, over all unique expressions in either of their programs, of the difference between the numbers of occurrences of the expression in the two programs. The number of species in a population was estimated from a sample of 64 organisms, using a “greedy” algorithm that attempts to merge the organisms into a minimum number of groups while ensuring that the maximum difference between organisms in a group does not exceed a specified species diameter (16 here, though we obtained similar results using other diameters). The degree of adaptation of a population was assessed from the fitness of the best organism in the population each generation.¹ We report the sum of the errors that the best organism produces over a set of inputs for the fitness-conferring problem (“fitness cases”); the lower this number is, the better the organism is adapted to its environment. We classified an entire run as “adaptive” if the average best fitness over the last 20 generations of the run was less than 200 (significantly better than is normally achieved by random individuals), and “non-adaptive” otherwise. We conducted 48 independent runs that included the sexual instructions and 48 independent runs without the sexual instructions.

Results and Analysis

Figure 1 shows the results of the 48 independent runs that included the sexual instructions. The graph plots *phases* of runs on axes of number of species (averaged over the phase) and average mother/child difference (also averaged over the phase). Each run begins with a reproductively incompetent phase; these phases are represented with square points. All of the runs in the study eventually achieved reproductive competence, and their reproductively competent phases (minus any generations at which there was a perfect solution to the fitness-

¹We did not use average fitness because large fitness penalties for arithmetic errors and nontermination, inspired by standard practice in GP, make such averages erratic.




Figure 1: Data from runs that included sexual operators.

conferring problem) are plotted using + for adaptive runs or × for non-adaptive runs. For runs that eventually produced perfect solutions to the fitness-conferring problem the phases during which they did so are plotted with diamond-shaped points.

The most obvious effect visible in the data is the clustering of most adaptive runs (in their reproductively competent, non-solution phases) between 50 and 64 (which is the maximum) diameter-16 species, a pattern clearly different from the results for non-adaptive runs, which span a range from about 5 to 64. Another notable effect is the clustering of the adaptive runs (in their reproductively competent, non-solution phases) between approximately 10 and 30 on the measure of mother/child differences. Non-adaptive runs exhibit a wider variance, ranging from less than 5 to approximately 35. This difference is significant with $p < 0.001$ using an F test.

Data from sexual vs. asexual runs appeared to differ only with respect to their reproductively incompetent phases, during which the asexual runs (not shown here) exhibited lower average mother/child differences and much lower numbers of diameter-16 species. This reflects the combinatorics of the available instruction set rather than any feature of an evolving population, as most organisms are random and there is no selection pressure during these phases. It is interesting that after reproductive competence the sexual and asexual runs are nearly indistinguishable. The data show that the sexual instructions are in fact used throughout the runs in which they are available, but their use does not seem to have an impact on mother/child differences or numbers of species.

A final set of observations concerns the effects of finding a perfect solution to the fitness-conferring problem. One might expect that the production (and maintenance) of a perfect solution would decrease the selective

advantage to species that diversify their children; when one is already doing as well as possible there would seem to be little advantage to be gained from further diversification. Indeed, in many runs the mother/child differences did decrease after a solution was found. But the effect seems to be more complex than this simple analysis would suggest. In some instances the mother/child differences *increased* after a solution was found and there appear to be effects on numbers of species as well. Figure 1 reveals an interesting pattern here; the diamonds (representing the solved phases) form a semicircle around the cluster of plusses (representing the unsolved phases of the adaptive runs). It appears that the production of a perfect solution is generally followed by movement away from the center of the cluster of points characterizing adaptive phases (and generally, but not always, in the direction of more species and/or smaller mother/child differences).

In summary the data show that the number of species in adaptive populations falls within a narrow, high range relative to non-adaptive populations. The data also show that the variance in the differences between mothers and their children is less for adaptive populations than for non-adaptive populations. In other words, in adaptive populations the species were more numerous and the diversification processes were more reliable. If the populations studied here were sub-populations within a larger environment then population-level selection effects (such as those called for in hierarchical theories of selection (Gould 2002)) would indirectly favor species-rich sub-populations with reliable diversification mechanisms. This provides one possible explanation for the evolution of diversifying reproductive systems.

Conclusions and Future Work

The reported data show patterns in the reproductive strategies of evolving populations of digital organisms. In the context of the studied parameters, adaptive populations were reliably diverse. In particular, the number of species in adaptive populations fell within a narrow, high range relative to non-adaptive populations, and the variance in the differences between mothers and their children was less for adaptive populations than for non-adaptive populations. These results are relevant to several open problems in artificial life, including challenges to achieve the transition to life in silico and to determine what is inevitable in the open-ended evolution of life.

The system used for the experiments, an “auto-constructive evolution” system called Pushpop, allows the reproductive and diversification mechanisms of the evolving organisms to evolve, thereby providing data not readily available from other digital organism systems. Pushpop runs produce large volumes of data and there are many other phenomena of biological interest that might be studied in this data. One avenue of research

that we are currently pursuing derives from the observation that Pushpop organisms can exhibit ontogenetic development (through the use of self-modifying code) and that this development, along with other aspects of their behavior, may vary depending on input and environment. This opens the possibility that Pushpop runs could be used to study relations between adaptation and developmental and behavioral plasticity, an issue of considerable current interest in evolutionary biology.

Acknowledgments Thanks are due to Rebecca S. Neimark, Mark Feinstein, Alan Robinson, Ray Copinger, and the Defense Advanced Research Projects Agency and Air Force Research Laboratory (agreement number F30502-00-2-0611).

References

- Edmonds, B. 1998. Meta-genetic programming: Co-evolving the operators of variation. CPM Report 98-32, Centre for Policy Modelling, Manchester Metropolitan University, UK, Aytoun St., Manchester, M1 3GH. UK.
- Gould, S. J. 2002. *The Structure of Evolutionary Theory*. Cambridge, MA: Harvard University Press.
- Koza, J. R. 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA: MIT Press.
- Maynard Smith, J., and Szathmary, E. 1999. *The Origins of Life*. Oxford University Press.
- Pargellis, A. N. 2001. Digital life behavior in the amoeba world. *Artificial Life* 7:63–75.
- Ray, T. S. 1991. Is it alive or is it GA. In Belew, R. K., and Booker, L. B., eds., *Proceedings of the Fourth International Conference on Genetic Algorithms*, 527–534. University of California - San Diego, La Jolla, CA, USA: Morgan Kaufmann.
- Spector, L., and Robinson, A. 2002. Genetic programming and autoconstructive evolution with the push programming language. *Genetic Programming and Evolvable Machines* 3(1):7–40.
- Spector, L. 2001. Autoconstructive evolution: Push, pushgp, and pushpop. In Spector, L.; Goodman, E.; Wu, A.; Langdon, W. B.; Voigt, H.-M.; Gen, M.; Sen, S.; Dorigo, M.; Pezeshk, S.; Garzon, M.; and Burke, E., eds., *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO-2001*, 137–146. Morgan Kaufmann Publishers.
- Stephens, C. R.; Olmedo, I. G.; Vargas, J. M.; and Waelbroeck, H. 1998. Self-adaptation in evolving systems. *Artificial Life* 4:183–201.
- Teller, A. 1996. Evolving programmers: The co-evolution of intelligent recombination operators. In Angeline, P. J., and Kinnear, Jr., K. E., eds., *Advances in Genetic Programming 2*. Cambridge, MA, USA: MIT Press. chapter 3, 45–68.
- Wilke, C. O.; Wang, J. L.; Ofria, C.; Lenski, R. E.; and Adami, C. 2001. Evolution of digital organisms at high mutation rates leads to survival of the flattest. *Nature* 412:331–333.
- Williams, G. C. 1966. *Adaptation and Natural Selection*. Princeton, NJ: Princeton University Press.