

Experiments in Digital Evolution (Editors' Introduction to the Special Issue)

Christoph Adami*
Claus O. Wilke
Digital Life Laboratory 136-93
California Institute of
Technology
Pasadena, CA 91125
{adami,wilke}@caltech.edu

Whether or not the field of artificial life has succeeded in doing what its name suggests, namely synthesizing life from non-living components, is a matter of contention. Clearly, this journal covers a broad range of topics related to the synthesis and simulation of living systems, but only a few articles go so far as to unabashedly study wholly artificial forms of life. The field of *digital evolution* is an exception: Artificial life forms, in the form of self-replicating computer code inhabiting specially prepared areas of a standard computer, have been used to learn about fundamental aspects of the evolutionary process since Tom Ray introduced us to them [23]. In this issue, we present experiments using digital organisms of the Avida variety (that is, implemented with the Avida software described in this issue [20]), but there are a number of other implementations of digital evolution that have been used for experimental evolution (see, e.g., [31, 32, 22, 5]). Whether or not these *digitals* are truly alive is ultimately of no concern to us as researchers: We use them because we are interested in complicated and vexing questions of evolutionary biology, and digitals offer us the possibility to attack them.

Digital evolution is currently undergoing a boom phase, and public perception of this discipline is steadily increasing [21]. This boom can be traced back in part to a maturation of the Avida software used in the majority of digital evolution experiments, in part to a perceived need for rigor in evolution experiments [7], and in part to the adoption of digitals as experimental organisms alongside bacteria and viruses by a growing number of microbiologists (see, e.g., [16]). Although work in digital evolution of the past decade has been reviewed by us recently [29], the versatility of the Avida software to conduct evolution experiments has never been displayed in the manner we have the opportunity to do in this special issue.

One of us (C.A.) has been teaching artificial life and evolution to advanced undergraduates and beginning graduate students at the California Institute of Technology since 1995 [1], and digital life has been a cornerstone of this class from the very beginning. The Avida software used in teaching this class, and developed expressly for research in evolutionary biology, was first written by C. Titus Brown and then by Charles Ofria in 1993 [19]. Since then, it has gone through many versions and revisions, with code contributions from a growing number of people, and a growing user base. But it is with the students taking CNS 175 (Artificial Life) and, since 2002, CNS 178 (Evolution and Biocomplexity) at Caltech that Avida has had its most lasting relationship. Each term, the students have to solve a number of problem sets, and at term's end, instead of a final examination, the students are asked to turn in a *final project* that uses Avida to perform an experiment in evolution. But while for CNS 175 the students could choose one of three carefully selected projects, for CNS 178 we decided to take a different approach. We would not only let the students answer a question, we would let them pose it, too. So, when final project time approached, the students were asked via a

* Present address: Keck Graduate Institute for Applied Life Sciences, 535 Watson Drive, Claremont, CA 91711.

“call for proposals”¹ to prepare a *research proposal*, which they subsequently had to defend in front of a committee [consisting of the teaching assistants (TAs)].

After evaluating the proposal (which contributed to the final grade), the TAs made recommendations intended to clarify the protocol and the question being addressed, and briefed the budding principal investigators on statistical hypothesis testing and error analysis. All final projects had to take the form of a research paper, with an introduction that placed the project in the proper context and the existing literature, a “Materials and Methods” section outlining the experimental techniques and materials used, a “Results” section, and a closing section with a discussion and conclusions².

The results of this experiment were spectacular: Of about 25 final projects handed in, we deemed that at least seven contained results that were publishable! Not only did the format of the final project force the students to think more like a federally funded scientist, it also made grading the homework much easier. After consulting with the editor of the journal *Artificial Life*, we extended an invitation to the authors of the seven best projects to expand or refine their work to make it publishable in a special issue, and five accepted. Their contributions are now collected here, along with an article that was a final project from CNS 175 in 2000, and which was published previously in the proceedings of the seventh artificial life conference [27]. Because it had not appeared in a journal and was closely related to the type of articles in this special issue, we decided to include it here. While the students were initially solely responsible for taking the data, analyzing it, and writing the report, they received considerable help from the editors as well as the TAs in the preparation of the final manuscripts, to an extent that it would seem unethical to leave the students as sole authors (as we originally intended). Also, the article by George Hagstrom ended up obtaining a considerable amount of input from colleagues at Michigan State University, where he spent a summer as a Caltech Summer Undergraduate Research Fellow (SURF). Incidentally, it is through the SURF program that Avida got its start, via an award to Titus Brown in the summer of 1993.

Opening the issue is an investigation into evolution in time-dependent (here, periodic) fitness landscapes [17]. The effect of non-static landscapes on evolution has been a perennial favorite to explain a number of puzzling aspects of evolving systems, but most of the work has been theoretical in nature, because performing experiments in changing landscapes is difficult and time-consuming (but see, e.g., [18] and [26]). But what is difficult or impossible with standard experimental organisms is just right for Avida. Initially, Ye Li wanted to study how the distance in time between two changes in a landscape would affect an adapting population. This project ultimately turned into an investigation of the dynamics of adaptation in a periodically changing fitness landscape (two peaks are periodically appearing and disappearing in an alternating manner). The study ultimately corroborates what is expected from quasi-species theory (the theory of self-replicating molecules under high mutation rates, also known as Eigen theory [6]), but also shows that there are subtle effects of adaptation that go beyond a simple mathematical description.

For the second contribution [28], Justin White decided to investigate an old paradigm of evolutionary biology: Is adaptation mainly driven by populations shifting from peak to peak (Wright’s shifting balance theory), or is the variation in a population’s fitness the determining factor (Fisher’s and Haldane’s view)? White came up with a clever scenario to test the importance of peak shifts: He constructed a landscape that could be inhabited by two very different populations (in other words, there were at least two peaks), but because it was a single-niche environment, both peaks could not be occupied at the same time. Thus, populations that were placed on a saddle point between the peaks

¹ The call can be found at <http://www.dllab.caltech.edu/cns178/handouts/call.pdf>.

² These formatting instructions can be viewed at the class Web site <http://www.dllab.caltech.edu/cns178/handouts/final-format.pdf>.

had to choose one or the other. This led to a bifurcation in evolutionary histories that was irreversible, and to the conclusion that peak shifts do seem to play an important role in evolutionary adaptation.

One of the least-studied areas in evolution is ecology (despite the fact that departments of “Evolution and Ecology” abound), because of the common belief that the time scales of evolution and ecology are seemingly so disparate. But now that experimental organisms with very short generation times are becoming more and more common, it has become increasingly important to consider evolution in an ecological context [3]. For his final project [15], Tyler Johnson wanted to investigate predator-prey cycles, the classic ecological paradigm, but instead of constructing digital predator and prey genomes, he decided to evolve pairs of organisms that prey on each other’s metabolic waste products. While the primary resource for all digitals is the energy provided to them via CPU cycles, an additional resource can boost replication speed if it can be utilized, that is, if an organism has the right genetic code to perform certain mathematical calculations that allow it to take up that resource. Johnson carefully set up the experiment in such a manner that the only two available resources were associated with different computations, with a fixed total amount (i.e., the resources did not naturally regenerate). Then, the experiment was started with one of the resources flush, up until a species evolved that could harvest it. By consuming it, this resource was turned into the second one. As the other resource became abundant, a second species could evolve that fed on it. After a while, oscillations between the two species, as well as between the resources, could be observed. While this scenario appears quite simple, the ensuing dynamics turned out to be much more complex than what the standard predator-prey equations would predict, and there was a distinct difference in the dynamics in the absence or presence of mutations, demonstrating once more that adaptation can play a significant role in ecological cycles.

Among the topics covered in the course were modern methods of phylogeny reconstruction [8], which prompted George Hagstrom (a freshman) to ask a nagging question that worries those who reconstruct phylogenies using standard algorithms. How good are these reconstructions [10]? Except for special cases [13], true phylogenies that we could test the reconstructions against do not exist, so the question has always seemed somewhat academic [12]. But digital life changed all that. Indeed, the first reconstructed digital phylogeny [14] appears to be that of the “stoned” computer virus (a virus that infected the boot sector of PCs between 1990 and 1995). In Avida, we can keep track of the main lineages in real time, and compare reconstructions against the actual sequence of events. While it is known that standard methods perform fairly well if evolution is mostly neutral, the impact of strong selection (for example, major extinction events in the history of a clade) are much more difficult to assess. Together with a group at Michigan State University that had investigated a similar question using digital organisms [11], Hagstrom studied the hypothesis that a reconstruction algorithm could only accurately group two organisms together if their direct ancestor had diverged from other groups *before* the last adaptive event that affected the clade.

The question that intrigued Jeffrey Edlund had to do with a phenomenon that garners more and more attention: Simple replicators adapt to a high mutation rate by reducing the deleterious effects of single mutations, or in other words, they develop robustness against mutations [30]. In biochemical life, adaptation to high mutation rates is common (a classical but extreme example is *Deinococcus radiodurans*), but there the mechanism of protection is usually well studied, and involves increasingly sophisticated DNA repair mechanisms [2]. Avidians do not have a mechanism to repair their genomes, so their ability to avoid mutational meltdown deserves some attention. Initially, Edlund proposed to study the importance of genetic *redundancy*, that is, the multiple occurrence of identical or similar genes within the genome, so that a knockout

of one does not destroy function. But a detailed analysis of this problem revealed that genetic redundancy could not play an important role, because even if it was present in one organism, it could not persist in the population in the absence of recombination. The reasoning goes as follows: If a genome with redundantly encoded function is hit by a mutation in a redundant gene, then its fitness is unchanged while at the same time the redundancy is irreversibly lost. Therefore, over time, any preexisting redundancy must fade from the population. And indeed, pervasive redundant coding could not be detected in mutationally robust organisms.

In a change of direction, we looked instead at the importance of *epistatic interactions* between loci in the genome and how they affect robustness [4]. Two genes are said to interact epistatically if the function of one gene depends on the function of the other. In some instances this interaction may be functionally necessary, but in other cases it may just be a case of bad coding, and the interdependence of modules could be reduced or even eliminated by rewriting the same genetic information in a more modular fashion. At high mutation rate, when the pressure to evolve mutational robustness is high, we would expect organisms to rearrange their genomes in order to remove epistatic interactions and increase their mutational robustness. This effect is subtle, but all evidence points to such robust recoding as the primary means of adaptation to high mutation rates for digital organisms.

The last research article in this issue originated from Daniel Wagenaar's final project in the class. While the other students whose articles are printed here were undergraduates when taking the class, Wagenaar was a graduate student in 2000, and his final project paper clearly showed this. With a little editing, it was accepted for the seventh artificial life conference, and it is included here (with some minor corrections and updating of references and links) because it shares with the other articles the same provenance—they originated as final term papers. The question Wagenaar set out to address is one of the thorniest of all evolutionary biology: What would we see if we replayed the course of evolution? The question was famously posed by Gould and Lewontin in the discourse about the "adaptationist programme" [9]. How important is historical contingency in shaping evolutionary outcomes? What is the influence of chance? What role does adaptation play? It seemed these questions would forever reside within the realm of speculation and theoretical modeling, but in a breakthrough development of the experimental method in evolutionary biology, Travisano and coworkers showed that it was possible to assess the relative importance of these factors in replicate populations of adapting bacteria [25]. Wagenaar set out to repeat these experiments using Avidians instead of *E. coli*, and showed that the effects seen earlier for biochemical organisms were very similar to what could be observed with digitals, but because of the great ease of doing experiments within the computer, he was able to provide better statistical accuracy, and could also follow, through time, the changing importance of chance, history, and adaptation in evolution.

Finally, we include in this issue a report giving an overview of the Avida digital life software [20]. The report describes such basics as the digital organisms themselves, the world they live in, and the environments they encounter, along with an introduction to the tools necessary to perform experiments with Avidians. The article is a guide to the usage of Avida, rather than a detailed reference manual. Nevertheless, it gives a good introduction to the more recent versions of Avida, for the first time since the original Avida manual was published in the class textbook, six years ago. This guide should be used in conjunction with the specific instructions and descriptions available as online documentation accompanying the current release of the software on its Web page at SourceForge.net [24] or at the Digital Life Lab's Web page.³ While the documentation

3 <http://www.dllab.caltech.edu/avida/v2.0/docs>.

is apt to change with increasing version numbers, we hope that most of the material covered by this guide will remain stable, and will not lose its relevance in the foreseeable future.

Acknowledgments

We are indebted to D. Allan Drummond and Daniel Wagenaar for their tireless efforts as teaching assistants for CNS 175 and 178, in particular for suggesting the idea of a proposal-based project system, and then making it happen. We also thank Charles Ofria for many discussions and for helping out students with technical and coding advice, and Richard Lenski for his unwavering commitment to digital life. We are grateful to Mark Bedau for giving us, and particularly the students whose first articles are published here, the opportunity to share our enthusiasm for digital evolution in this special issue. Finally, we thank the referees who have kindly agreed to turn these articles around in record time, and delivered. This work was supported by the National Science Foundation's Biocomplexity Program under Contract No. DEB-9981397.

References

1. Adami, C. (1998). *Introduction to artificial life*. New York: Springer-Verlag.
2. Battista, J. R., Earl, A. M., & Park, M.-J. (1999). Why is *Deinococcus radiodurans* so resistant to ionizing radiation? *Trends in Microbiology*, *7*, 362–365.
3. Bohannan, B. J. M., & Lenski, R. E. (2000). Linking genetic change to community evolution: Insights from studies of bacteria and bacteriophage. *Ecology Letters*, *3*, 362–377.
4. Edlund, J. A., & Adami, C. (2004). Evolution of robustness in digital organisms. *Artificial Life*, *10*, 167–179.
5. Egri-Nagy, A., & Nehaniv, C. L. (2003). Evolvability of the genotype-phenotype relation in populations of self-replicating digital organisms in a tierra-like system. *Lecture Notes in Artificial Intelligence*, *2801*, 238–247.
6. Eigen, M. (1971). Self-organization of matter and the evolution of macromolecules. *Naturwissenschaften*, *58*, 465–523.
7. Elena, S. F., & Lenski, R. E. (2003). Evolution experiments with microorganisms: The dynamics and genetic bases of adaptation. *Nature Reviews Genetics*, *4*, 457–469.
8. Felsenstein, J. (2004). *Inferring phylogenies*. Sunderland, MA: Sinauer Associates.
9. Gould, S. J., & Lewontin, R. C. (1979). The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist programme. *Proceedings of the Royal Society of London Series B*, *205*, 581–598.
10. Hagstrom, G. I., Hang, D. H., Ofria, C., & Torng, E. (2004). Using Avida to test the effects of natural selection on phylogenetic reconstruction methods. *Artificial Life*, *10*, 157–166.
11. Hang, D., Ofria, C., Schmidt, T., & Torng, E. (2003). The effect of natural selection on phylogeny reconstruction algorithms. *Lecture Notes in Computer Science*, *2723*, 13–24.
12. Hillis, D. M. (1995). Approaches for assessing phylogenetic accuracy. *Systematic Biology*, *44*, 3–16.
13. Hillis, D. M., Bull, J. J., White, M. E., Badgett, M., & Molineux, I. J. (1992). Experimental generation of a known phylogeny. *Science*, *255*, 589–591.
14. Hull, D. B. (1995). Naming and classifying computer viruses. <http://minerva.pom.devry.edu/~dhull/vclass1.htm>.
15. Johnson, T. J., & Wilke, C. O. (2004). Evolution of resource competition between mutually dependent digital organisms. *Artificial Life*, *10*, 145–156.
16. Lenski, R. E. (2001). Twice as natural. *Nature*, *414*, 255.

17. Li, Y., & Wilke, C. O. (2004). Digital evolution in time-dependent fitness landscapes. *Artificial Life*, *10*, 123–134.
18. Novella, I. S., Hershey, C. L., Escarmis, C., Domingo, E., & Holland, J. J. (1999). Lack of evolutionary stasis during alternating replication of an arbovirus in insect and mammalian cells. *Journal of Molecular Biology*, *287*, 459–465.
19. Ofria, C., Brown, C. T., & Adami, C. (1998). The Avida user's manual. In [1].
20. Ofria, C., & Wilke, C. O. (2004). Avida: A software platform for research in computational evolutionary biology. *Artificial Life*, *10*, 191–229.
21. O'Neill, B. (2003). Digital evolution. *PLoS Biology*, *1*, 11–14.
22. Pargellis, A. N. (2001). Digital life behavior in the Amoeba world. *Artificial Life*, *7*, 63–75.
23. Ray, T. S. (1992). An approach to the synthesis of life. In C. Langton, C. Taylor, J. Farmer, & S. Rasmussen (Eds.), *Proceedings of Artificial Life II*. Reading, MA: Addison-Wesley.
24. SourceForge.net. <http://sourceforge.net/projects/avida>.
25. Travisano, M., Mongold, J., Bennett, A., & Lenski, R. (1995). Experimental tests of the roles of adaptation, chance, and history in evolution. *Science*, *267*, 87–90.
26. Turner, P. E., & Elena, S. F. (2000). Cost of host radiation in an RNA virus. *Genetics*, *156*, 1465–1470.
27. Wagenaar, D. A., & Adami, C. (2000). Influence of chance, history, and adaptation on evolution of Digitalia. In M. Bedau, J. McCaskill, N. Packard, & S. Rasmussen (Eds.), *Proceedings of Artificial Life VII* (pp. 216–220). Cambridge, MA: MIT Press.
28. White, J. S., & Adami, C. (2004). Bifurcation into functional niches in adaptation. *Artificial Life*, *10*, 135–144.
29. Wilke, C. O., & Adami, C. (2002). The biology of digital organisms. *Trends in Ecology & Evolution*, *17*, 528–532.
30. Wilke, C. O., & Adami, C. (2003). Evolution of mutational robustness. *Mutation Research/Fundamental and Molecular Mechanisms of Mutagenesis*, *522*, 3–11.
31. Yedid, G., & Bell, G. (2001). Microevolution in an electronic microcosm. *American Naturalist*, *157*, 465–487.
32. Yedid, G., & Bell, G. (2002). Macroevolution simulated with autonomously replicating computer programs. *Nature*, *420*, 810–812.