Book Review: "Introduction to Artificial Life"
Christoph Adami, Springer-Verlag, New York, 1998

Tim Taylor
International Centre for Computer Games and Virtual Entertainment
University of Abertay Dundee
Dundee, DD1 1HG, Scotland
t.taylor@tay.ac.uk

Despite the popularity of artificial life over the last 15 years or so, as witnessed by the establishment of several international conferences on the subject and the birth of the *Artificial Life Journal* and similar publications, there has been a conspicuous lack of decent introductory texts on the subject. Notable exceptions include "*Artificial Life: An Introduction*" (Langton 1995), an eclectic selection of papers covering a wide range of subjects that fall under the "artificial life" umbrella, and "*The Philosophy of Artificial Life*" (Boden 1996), another collection of papers, this time dealing with the philosophical underpinnings and implications of the subject. A handful of popular science books also introduce the subject, Levy's arguably still being the best (Levy 1993), but someone looking for a textbook suitable for an undergraduate or graduate course on artificial life may be disappointed.

Christoph Adami's "*Introduction to Artificial Life*" is a welcome attempt to help fill this gap. The text has grown out of lectures given to advanced undergraduate- and graduate-level students in Computation and Neural Systems and in Physics at Caltech (USA) since 1995. The content will also be of value to professional researchers with an interest in the analysis of artificial evolutionary systems.

The first point to note about the text is that, despite the title, it is emphatically *not* a general introduction to the field of artificial life. Instead, it concentrates very specifically on introducing statistical mechanical and information theoretic techniques for analyzing evolutionary processes. In Chapter 1, Adami does give a brief introduction to the broader field, including evolution and learning in physically modeled virtual creatures, biologically inspired robotics, emergent behavior in multi-agent systems, in vitro evolution of RNA molecules, and cellular automata.

The remainder of the book is largely devoted to the introduction of a battery of theoretical approaches from the physical sciences and a discussion of their application to the analysis of evolutionary systems (both real and artificial). Subjects covered include an introduction to Information Theory, statistical mechanics and thermodynamics, informational definitions of complexity, self-organized criticality, percolation theory, fitness landscapes, the propagation of information, and adaptive learning at the error threshold. Background is assumed in fundamental methods of statistical physics, knowledge of basic computer architecture and programming skills, as well as rudimentary knowledge of biology. Many mathematical results are presented without derivation, although suggestions for further reading are generally given in such cases.

Throughout the book, the analysis tools developed are, from time to time, applied to Adami and colleagues' **avida** system. The model, derived from Ray's well-known **tierra** (Ray 1991), simulates the parallel execution and evolution of a population of self-replicating computer programs. The design of **avida** is not discussed at length in the main text, but rather in the appendix, where the system's User Manual is reprinted. A CD ROM is supplied with the book, which contains versions of **avida** for Windows 95, Windows NT and Unix systems (but not Macintosh). The CD ROM also contains other supporting material, such as source code for programs to demonstrate aspects of self-organized criticality and percolation theory, as well as an HTML document containing links to sources of further reading on many of the subjects discussed in the book. The latter is a particularly good idea, although, perhaps inevitably, some of the links are now broken (the book, together with the CD, was published in 1998).

While the work introduced is doubtlessly valuable and provides new insights into the dynamics of populations of self-replicating strings, I found myself asking how it fits in with the characterization of artificial life provided by Chris Langton (the man who was largely responsible for rekindling interest in the field in the mid-1980s). Langton describes artificial life as:

'[the study] of man-made systems that exhibit behaviors characteristic of natural living systems. It complements the traditional biological sciences concerned with the *analysis* of living organisms by attempting to *synthesize* life-like behaviors within computers and other artificial media. By extending the empirical foundation upon which biology is based beyond the carbon-chain life that has evolved on Earth, Artificial Life can contribute to theoretical biology by locating *life-as-we-know-it* within the larger picture of *life-as-it-could-be*.' ' (Langton 1988)

In other words, artificial life is distinguished from traditional theoretical biology by its fundamentally synthetic, rather than analytic, approach. I believe that most people would broadly accept this distinction. Although *Introduction to Artificial Life* describes occasional experiments on **avida**, the emphasis of the book is clearly on the development of a variety of analytical tools for understanding evolution (both natural and artificial). Of course, the behavior of artificial life simulations needs to be analyzed just as carefully as any other kind of scientific model, and Adami's work represents a valuable and much-needed contribution to such analysis. However, nearly all coverage of the *design* of **avida** is confined to the appendix, and even here it is described with little discussion of the theory behind its design. This I consider a great pity, as it is precisely in the process of designing and building such systems that one is faced with a myriad of choices and design decisions, even the smallest of which may turn out to radically alter the resulting behavior of the system. It is in highlighting new questions about the design and construction of systems – questions that simply do not arise when analyzing the behavior of a predefined, given system – that artificial life perhaps has the greatest potential to complement the results of analytical models and thereby contribute to a broader understanding of life.

When Adami does turn to more philosophical considerations of the nature of life, it is clear that he regards information and computation as the critical issues. For example, when discussing the self-replicating programs in **tierra** and **avida**, he suggests "because of the remaining computational-universality of their virtual CPUs … we feel compelled to speculate that such constructions constitute *minimal living systems*" (p.58, original emphasis). To motivate the discussion of information theory, Adami asserts "the discovery of the genetic code cemented the fact that information is the *central* pillar in any attempt to understand life" (p.59, original emphasis).

Despite an observation in Chapter 1 that "life seems to be a property of a *collection* of components but not a property of the components themselves" (p.4, original emphasis), which suggests that the *organization* of organisms is a (or, some would argue, *the*) defining feature of life, Adami does not expand upon this theme. Throughout the rest of the book he seems happy to regard self-replicating strings (be they computer programs or biomolecules) as "minimal living systems". This lack of consideration of organizational issues is also reflected in the approach to modeling populations of organisms; much of the theory is based on assumptions like the population comprises only a single species (e.g. Chapter 4), or in other cases, that ecological interactions and coevolution between species are negligible (e.g. Chapter 8). Of course, such assumptions are introduced to make the analysis tractable, but I suspect that it is precisely in these questions of organization, both at the level of the individual organism, and at the level of the ecosystem, that living systems may be most appropriately distinguished from non-living. For example, the most interesting (most lifelike?) results to arise from **tierra** concerned the emergence of ecological relationships, such as parasitism, between the evolving programs (Ray 1991). Models and analyses that concentrate on individual self-replicators evolving in static environments will fail to elucidate the important influences of such ecological interactions.

Notwithstanding the foregoing remarks, the analysis techniques introduced by Adami are clearly relevant in certain situations. The book is interesting and informative, and represents a valuable contribution to the literature. It will be of great interest to any researcher looking for new ways to analyze and interpret the behavior of an artificial evolutionary system. My main gripe is with the title; for someone looking for an introductory text on artificial life in its full breadth, I would still recommend (Langton 1995) and (Boden 1996) as initial sources, perhaps reading Adami's more specific text on the analysis of evolutionary systems as a secondary text, or progressing to it (if this is where one's particular interests lie) at a later stage.

References:

(Boden 1996) Margaret A. Boden, editor, *The Philosophy of Artificial Life.* Oxford University Press, 1996.

(Langton 1988) Christopher G. Langton. Artificial Life. In C.G. Langton, editor, *Artificial Life*, volume VI of *Santa Fe Institute Studies in the Sciences of Complexity*, pp.1-47, Addison-Wesley, Reading, MA, 1988.

(Langton 1995) Christopher G. Langton, editor, *Artificial Life: An Introduction.* MIT Press, 1995.

(Levy 1993) Steven Levy. *Artificial Life: A Report from the Frontier Where Computers Meet Biolog*y. Vintage Books, 1993.

(Ray 1991) Thomas S. Ray. An approach to the synthesis of life. In C.G. Langton, C. Taylor, J.D. Farmer and S. Rasmussen, editors, Artificial Life II, pp. 371-408. Addison-Wesley, Redwood City, CA, 1991.