Artificial Life: Discipline or Method?

Report on a Debate Held at ECAL'99

Jason Noble*, Seth Bullock† and Ezequiel A. Di Paolo‡

*Center for Adaptive Behavior and Cognition, MPI für Bildungsforschung, Berlin

†School of Computer Studies, University of Leeds

‡GMD—German National Research Center for Information Technology (AiS)

noble@mpib-berlin.mpg.de, seth@scs.leeds.ac.uk

Ezequiel.Di-Paolo@mail.gmd.de

How can artificial life (AL) advance scientific understanding? Is AL best seen as a new discipline, or as a collection of novel computational methods that can be applied to old problems? And given that the products of AL research range from abstract existence proofs to working robots to detailed simulation models, are there standards of quality or usefulness that can be applied across the whole field? On September 16th 1999 in Lausanne, Switzerland, a debate on these questions was held as part of the Fifth European Conference on Artificial Life. As the organizers, we wanted to foster a constructive discussion regarding the scientific status, and future, of AL. We were well aware that some of these issues had been raised before (e.g., Miller, 1995) but we felt that earlier treatments had perhaps not reached a wide enough audience. The format for the debate consisted of contributions from invited panelists followed by an open discussion. The panelists were Chris Langton, Mark Bedau, Simon Kirby, and Inman Harvey—Hiroaki Kitano was scheduled to participate but regrettably could not attend the conference.

We started by sketching a continuum of approaches to AL as science: at one extreme, there are researchers who use techniques such as genetic algorithms and animat-style simulations to look at existing problems. Typically the problems come from within biology. The work of Kitano and his colleagues (e.g., 1997) on morphogenesis in *Drosophila* is one example. We see this work as exemplifying the idea of AL as a method, or a collection of methods, that could (at least in theory) be put to use by investigators in many different fields. At the other extreme is the view that AL opens up whole new ways of thinking; that it is a discipline in itself. An example of this is Ray's 1994 work on Tierra, that Ray and some other commentators have taken to raise fundamental issues about what it means to call a system alive.

If the existence of the continuum is granted, two questions are raised. First, are all of the possible positions along it tenable? The skeptic might ask whether computer simulations of the kind developed by AL researchers ever add anything to existing formal methods in a discipline like biology. At the opposite end of the spectrum, one might be cynical about the possibility of objectively studying "life as it could be." The second question concerns quality: how can we distinguish good work from bad? The two questions are not independent: if one sees AL research as some kind of thought experiment, one's quality criteria may well differ from those of someone who is interested in more-or-less precise models of real-world systems.

These were the terms within which we framed the debate. We then encouraged each panelist to present his own views on the field, and requested that the audience return to questions about the scope of AL and the problem of quality in the open discussion.

Chris Langton (The Swarm Corporation) began by offering definitions of terms such as "method," "discipline," and "tool" that had been used in the debate abstract.[1] He argued that methods and disciplines typically overlap, and thus deciding whether AL is one or the other is a moot point. The key question is whether or not there exist research topics unique to AL. Langton answered this in the affirmative: AL's unique content is the study of synthesis. He noted that synthetic methods had originated with the pioneers of computing (Turing and von Neumann) but had not been taken up until many years later. Synthesis allows the investigator to go beyond the phenomena that happen to exist in Nature, and to create new phenomena. Langton argued that theoretical biology, for example, should be about the organisms and ecosystems that might evolve if we could "replay the tape." Thus AL gives us access to the possible as well as the actual, allowing us to discover laws of greater generality.

Mark Bedau (Dept. of Philosophy, Reed College, Oregon) described his interest in the fundamental properties of living systems. He argued that one profitable use of AL techniques is to allow us to find the simplest system capable of reproducing some key property of life, such as hierarchical emergence or unbounded complexity. Discovering such a minimal system would have obvious benefits for our understanding of the property in question. Bedau felt that AL, in this mode at least, deals in computational thought experiments rather than realistic simulations. He claimed that the way to keep these models grounded in reality is to develop operational metrics for the phenomena of interest, and, using these metrics, to compare artificial systems to their natural analogues. Bedau concluded by noting that the long-run success of his approach is an empirical question that turns on facts about how the universe is organized: will simple general principles allow us to explain many natural systems, or is the devil in the details?

Simon Kirby (Dept. of Linguistics, University of Edinburgh) pointed out that, although the majority of extant work in AL is related to biology, the

[1] For the original abstract and other materials associated with the debate, refer to `http://www-abc.mpib-berlin.mpg.de/users/noble/ECALDebate/`

perspective that AL methods afford can illuminate problems from many fields. Kirby pointed out that the study of language has involved explicit but incompatible idealizations about what language is. For example, Chomsky's idealization of language as the knowledge of a single speaker in a homogeneous speech community at a particular point in time contrasts with the historical idealization of language as a community-wide phenomenon subject to static laws of change. Kirby expressed the hope that the special perspective provided by AL might unify these competing idealizations by modelling explicitly how local interactions can give rise to global phenomena. In this way AL can throw new light on disciplines outside biology.

Inman Harvey (Centre for Computational Neuroscience and Robotics, University of Sussex) expressed the opinion that AL is not a discipline in itself, nor should it be. He pointed out that crossover between disciplines leading to "flaky speculations and hot air" may be one of the great benefits of AL as a movement. Harvey cautioned that AL should not expect overnight acceptance by other sciences. He mentioned the series of workshops entitled "Towards a Theoretical Biology" organized by C. H. Waddington in the late 1960s—today, the contents of these workshops would be recognized as AL, tackling the same topics, although without the aid of powerful computers. It has taken time for the methods and questions of theoretical biology to become an accepted part of biology. Similarly, we can expect that it will take a period of decades before the methods and questions of contemporary AL will be widely accepted. The best way to encourage this acceptance is to work directly with biologists. Harvey concluded by stressing that, while pursuing this goal, AL researchers should continue to explore crossover between disciplines and generate the "flaky stuff" that such interdisciplinary interaction often results in.

Several themes were touched upon in the ensuing open discussion. Pleasingly one of these was the question of how to judge whether AL research was of high quality. Three clear types of answer were put forward.

Mark Bedau expanded on his argument that well-defined operational metrics, sensitive to the phenomena claimed to be central to AL, must be constructed. He gave the example of open-ended evolution. How are we to measure the potential for open-ended evolution in an artificial system such as Tierra? How are we to compare this measure to that of a natural system such as the fossil record? Once a metric is established, researchers will be in a position to assess the ability of models to account for naturally observed phenomena and to construct the simplest model possible that exhibits these phenomena.

Chris Langton raised the notion of a model's ability to allow new predictions to be made and subsequently tested in the real world. He gave the example of the power law distribution of extinction events manifested by coevolutionary models and the subsequent prediction that natural extinction events should obey a similar law.

Takashi Ikegami (University of Tokyo) suggested that rather than merely providing predictions of what must be the case in living systems, i.e., describing the unifying features of "life as it could be," AL research should also attempt to provide negative hypotheses: "life as it cannot be." With claims such as these in hand, biologists are in a position to falsify the AL models that gave rise to them, and hence move AL forward.

Much of the debate focussed on whether AL systems were best thought of as models, experiments, automated thought experiments, intuition pumps, or conceptual frameworks, and what influence this might have on how to design and assess such systems. Whilst the field as a whole appears to be far from a consensus on this issue, the notion that a plurality of approaches is healthy seems widespread. Inman Harvey suggested that the most abstract AL models might be best construed as attempts to build new calculi, new tools with which to address problems from theoretical biology or elsewhere.

Issues that remained unresolved, and that are perhaps ripe for further debate, include how to ensure that the status of an AL system (automated thought experiment vs. nascent calculus, for instance) is clearly understood and conveyed, and how the status of an AL system has implications for what conclusions can be legitimately drawn from the system.

An issue upon which most of the discussants were in agreement concerned the scope of AL's subject matter. The panelists and audience agreed with Simon Kirby's assertion that AL, like complex systems theory, had a positive contribution to offer to many disciplines, including philosophy, linguistics, economics, psychology, geography, archaeology and even theology.

There was less agreement as to whether there were questions central to AL that were not also the concern of other disciplines, i.e., a core of research topics unique to AL. Whilst some contributors saw all AL as ultimately taking place within and augmenting some established research paradigm, typically theoretical biology, others saw the potential for AL to reveal unifying principles that were effective across disciplines that were otherwise unrelated—thus, AL might reveal fundamental similarities between economies and ecologies. Other contributors allowed that the development and assessment of AL tools and techniques might be considered to be central to AL in much the same way that the development and assessment of statistical tools by statisticians is to some extent divorced from their application in fields such as psychology or physics.

Questions of AL's application in other disciplines immediately raised what was perceived to be a serious problem: how can AL perspectives, models and techniques come to gain acceptance in these other fields? Titus Brown (Caltech) related that finding collaborators from within relevant research fields may take time, but that such collaboration can result in work with added credibility. Inman Harvey stressed that conservative models which only minimally perturb orthodox paradigms may have the greatest chance of success; whether this is the only way forward is perhaps a topic for future debate. What appears clear is that AL as a science is seeking ways of interfacing and interacting with other sciences, and that there are interesting and challenging problems to be faced in the course of this enterprise.

In closing, we would like to raise two issues which emerged from the debate for further discussion, perhaps in a similar debate at some future AL conference.

First, given that it was generally felt that the success of AL research as science might be measured by its impact on other disciplines, what do scientists from other disciplines want from AL? We suspect that these scientists do not just want to be shown examples of phenomena similar to their own instantiated on computers. They want answers to questions which they find hard to address with the techniques at their disposal; answers to questions such as why is human language the only signalling system yet documented to exhibit recursive, compositional syntax? Why do the extinction events evidenced by the fossil record exhibit the particular pattern that they do? It would be advantageous to have a way of dividing such questions into those amenable to an AL approach and those not so amenable. Must we merely attempt each project and see what happens, or can we offer principled criteria for the successful application of AL methods?

The second issue poses a related question: what kind of answers can AL research provide? A typical characterisation of AL is as a source of new data points; a way of widening an empirical database by providing synthetic examples of the natural phenomenon of interest. When augmented by data from AL, a discipline can separate the regularities which underlie all possible cases of a phenomenon from the idiosyncrasies of the particular natural examples we happen to have access to. However, many of the discussants at this debate do *not* see their work as a source of empirical data. They see it as a source of new ways of thinking, or novel intuitions; as a way of testing the coherence of theories, or of generating other non-empirical results. Is it time to forgo the philosophically rather difficult notion that the systems developed by AL researchers have the same status as the natural phenomena they seek to resemble? We believe that these artificial systems are more like conventional models, built in order to clarify and extend theories of natural phenomena, rather than to augment natural phenomena with artificial brethren.

References

Kitano, H., Hamahashi, S., Kitazawa, J., Takao, K., & Imai, S.-i. (1997). The virtual biology laboratories: A new approach of computational biology. In P. Husbands & I. Harvey (eds.), *Proceedings of the Fourth European Conference on Artificial Life (ECAL'97)*, pp. 274–283, MIT Press / Bradford Books, Cambridge, MA.

Miller, G. F. (1995). Artificial life as theoretical biology: How to do real science with computer simulation. Cognitive Science Research Paper 378, School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK.

Ray, T. S. (1994). An evolutionary approach to synthetic biology: Zen and the art of creating life. *Artificial Life*, *1*(1/2), 179–209.