

Artificial Life as Theoretical Biology:

How to do real science with computer simulation

Geoffrey F. Miller

Cognitive and Computing Sciences

University of Sussex

Falmer, Brighton BN1 9QH, England

geoffm@cogs.susx.ac.uk

Abstract

Artificial Life (A-Life) research offers, among other things, a new style of computer simulation for understanding biological systems and processes. But most current A-Life work does not show enough methodological sophistication to count as good theoretical biology. As a first step towards developing a stronger methodology for A-Life, this paper (1) identifies some methodological pitfalls arising from the ‘computer science influence’ in A-Life, (2) suggests some methodological heuristics for A-Life as theoretical biology, (3) notes the strengths of A-Life methods versus previous research methods in biology, (4) examines some open questions in theoretical biology that may benefit from A-Life simulation, and (5) argues that the debate over ‘Strong A-Life’ is not relevant to A-Life’s utility for theoretical biology.

1 Introduction: Simulating our way into the Dark Continent

Complex biological adaptation is the dark continent of science. We have recognized it and traced its boundaries, but its verdant, savage interior resists our conceptual cartography and eats scientists alive. Few who venture to explore it ever return to the comforts of reductionism and creationism.

February 1, 1995

2

This romantic, colonialist view seems implicit in much of Artificial Life (A-Life) research. It maintains a kind of ‘Orientalism of the Organic’ that views complex adaptive systems as newly discovered things and evolution, development, and learning as newly discovered processes. For those who come to biology from computer science, this attitude may be accurate in a strictly autobiographical sense: many A-Life researchers learned to program computers before they learned to look at organic complexity. But like all Orientalism (Said, 1978), the Orientalism of the Organic ignores, exploits, or patronizes the indigenous natives who have flourished on the Dark Continent for generations, and who have mapped its topography in their own language. I refer, of course, to evolutionary biologists, who have made their home in the study of complex biological adaptation ever since Darwin.

This paper examines A-Life as theoretical biology, as a set of computer simulation methods that may prove useful to biologists given their native concerns. I will not address A-Life as engineering, entertainment, pedagogy, philosophy of biology, or runaway post-modern cult. My aims are methodological rather than metaphysical. In Sober’s (1989) terms, I am concerned with ‘weak A-Life’ (computer programs as useful simulations of real life) rather than ‘strong A-Life’ (computer programs as actually living in some respects). Just as philosophical debates about ‘strong AI’ distracted attention from serious methodological issues in using computers to model psychological processes, I am worried that debates about strong A-Life will lead to a neglect of methodology in A-Life.

My arguments may undermine some of A-Life’s philosophical pretenses by viewing most A-Life research as theoretical biology done badly ... so far. These arguments will offer little solace to those who view A-Life as the cybernetic eco-religion of the future. All I can offer is the hope that, *if successful*, A-Life will be absorbed entirely into evolutionary biology as a useful new modelling tool. This will seem no poor fate to those of us who view evolutionary biology as the most fascinating and

important body of knowledge ever developed by our species.

Taylor and Jefferson (1994) have also addressed A-Life as a tool for biological inquiry, pointing out that A-Life methods are "powerful enough to capture much of the complexity of living systems, yet in a form that is more easily manipulable, repeatable, and subject to precisely controlled experiment than are the corresponding natural processes." They reviewed the diversity of A-Life work, including wetware, software, and hardware, and compared mathematical versus simulation methods for understanding evolution. They also listed some open problems in biology that are amenable to A-Life modeling, including the origin of life, self-organization, cultural evolution, the origin of sex, Wright's (1932) "shifting balance" theory, the relation between fitness and adaptedness, the structure of ecosystems, and the nature of mind. This paper complements theirs by taking a more critical attitude towards the 'computer science influence' in A-Life, by focusing on methodological problems, and by developing a more inclusive list of open biological problems. Other methodological papers concerned at least somewhat with A-Life as theoretical biology include Bedau & Packard (1992), Cariani (1992), Cliff (1991), Godfrey-Smith (1994), Langton (1989), Pattee (1989), and Sober (1992).

2 Computer Science versus Real Science

A-Life developed historically as an outgrowth of computer science, and many A-Life researchers were trained originally as computer scientists. This situation has created problems for A-Life as a branch of biological science, because most computer science (CS) differs from real science (RS) in several relevant ways. Fundamentally, RS analyzes pre-existing natural phenomena such as stars, proteins, animal behaviors, or economies. By contrast, even the purest forms of CS can claim only to investigate the mathematical and logical foundations of 'computation' – something that occurs only at a certain level of description in specially-constructed human artefacts that are themselves

February 1, 1995

4

designed through CS research. CS methods are much closer to analytical methods used by engineers (such as finite element analysis, stability analysis, or perturbation analysis) than to empirical methods used by scientists. This basic distinction leads to several other differences between RS and CS, which I will list with brazen simplicity (and tongue slightly in cheek).

RS seeks knowledge of nature, whereas CS seeks profit through technical improvements in hardware and software. RS chooses problems for their theoretical and practical importance, whereas CS chooses problems for their economic relevance. RS analyzes existing natural systems through developing speculative theories subject to hypothesis-testing by observation, experimentation, simulation, whereas CS builds new artificial systems through engineering and debugging. RS requires skills in observation, experimentation, comparative analysis, statistics, scholarship, interdisciplinary communication, and bold imagination, whereas CS requires skills in programming and mathematical analysis. RS rewards discoveries for their theoretical generality, experimental clarity, and practical applicability, whereas CS rewards inventions for their patentability, marketability, and complexity. RS advances through public, loosely organized collaborations (called "research areas") by huge numbers of researchers working over many years, whereas CS advances through largely private, more tightly organized collaborations (called "project teams") by smaller numbers of programmers working over several months. RS justifies public research funding by appeals to intellectual curiosity and human welfare, whereas CS justifies private funding as a necessary business investment in product research and development.

Perhaps I am making unfair comparisons between pure natural science and commercial software development. All sciences range from pure theory through experiment, observation, analysis to applications and engineering; all sciences are subject to social, political, and economic pressures. But the key difference remains: the phenomena studied by natural science predate the science itself,

February 1, 1995

5

whereas the phenomena studied by computer science (e.g. ‘computation’) depend on the science for their very existence. Computer science is more similar in nature and spirit to architecture and aeronautical engineering than it is to physics or biology.

Granted, computer science has been spectacularly successful as an engineering discipline: computer speeds and computer sales have grown exponentially over several decades, continuing to double every few years. But computer scientists run into trouble when they try to do real science, because they are simply not trained for it. Whenever computer scientists try to do real psychology for example, a boom-and-bust cycle results, as in artificial intelligence during the 1960s, cognitive science during the 1970s, and neural networks research during the 1980s. With suggestive initial results come wild promises of further progress and massive influxes of research funding; intellectual stagnation sets in as the promised conceptual break-throughs remain elusive, but modestly useful real-world applications keep the field limping along for another few years. I am afraid that A-Life will repeat this same pattern.

Again and again, the same problems arise when computer scientists develop a new field that claims to be real science. *Poor scholarship*: the field shows historical amnesia, interdisciplinary blindness, and ignorance of current work in the relevant existing sciences. *Poor research methods*: the field lacks explicit hypothesis-testing, systematic observation, controlled experimentation, and statistical analysis. *Poor analysis of results*: the field avoids recognizing or understanding its failures and overgeneralizes its successes in wildly inappropriate ways. *Poor follow-through*: the field jinks from one research fad to another, failing to replicate and extend its findings in ways that could lead to a conceptually integrated discipline. Many of these problems result from an ‘engineering mentality’ that seeks to build impressively complicated ‘masterpiece’ systems rather than to develop simple theories that explain complex phenomena. Masterpieces of computer programming (such as those developed by

Ph.D. students in artificial intelligence or cognitive science) are almost always useless as scientific models because they usually require too much pre-processing of the input, too generous an interpretation of the output, and too many hidden assumptions and ‘hacks’.

These precedents suggest that unless A-Life is very careful, it will become a historical curiosity in short order – after being more or less successful in squandering many millions of dollars of public research funding that could have been better spent on a few bright evolutionary biologists and evolutionary psychologists. The remainder of this paper suggests some ways we can learn from the mistakes of the past, by identifying some methodological heuristics for doing A-Life as good theoretical biology, and some areas of biology that may benefit especially from A-Life simulations.

3 Six Methodological Heuristics for A-Life

1. *Identify a known, unsolved problem in theoretical biology that could be addressed using simulation*

This step is much harder than it sounds. Biology is a mature, successful science that has become quite sophisticated over the last few decades. One cannot just pick up a copy of *The Selfish Gene* (Dawkins, 1976) or an introductory undergraduate biology textbook, find an interesting-sounding issue, and forge ahead with a simulation. A basic maxim of modern science is: *if you read about some intriguing topic through a popular science book or textbook, there have probably been at least a hundred papers published on it, and most of the interesting issues have already been resolved.* A second maxim, more specific to theoretical biology, is that if the interesting issue hasn’t yet been resolved in the literature, it will be resolved just before you finish programming your A-Life simulation, probably by some graduate student at Oxford, in a maddeningly brilliant proof that you won’t understand anyway. Even worse, students in theoretical biology are beginning to use computer simulation routinely as a supplement to formal mathematical analysis, and they know the relevant

biological issues much better than A-Life researchers do.

So, it is difficult to find a significant unsolved problem that can be addressed through A-Life computer simulations, and that has not already been addressed by standard methods of theoretical biology (e.g. verbal argument, formal population genetics models, optimality theory, evolutionary game theory, or dynamical programming). Most such problems can probably be found (1) on the cutting edge of some hot area in theoretical biology, (2) at the intersection of two current areas, or (3) in the dustbin of history, as an issue raised decades ago (probably by Darwin, Baldwin, Fisher, or Haldane) and then forgotten. In any case, it is almost impossible to choose an appropriate research question without a thorough knowledge of the biological literature. Section 5 identifies a few areas of theoretical biology that seem open to timely A-Life work.

2. Collaborate with real biologists who have already worked on the problem.

There are only two ways to do well-informed theoretical biology: become an expert, or talk with an expert. The second way is much easier. Since biologists disagree with one another about the open questions in biology, it is best to consult with several, broadening one's appreciation of the issues.

Theoretical biologists are an ornery, bristly lot, and take great pride in their mathematical skills and knowledge of biology. They should be approached with respect. One should avoid suggesting that A-Life simulations will solve all the problems of biology in the next ten years, and they had better jump on the bandwagon if they know what's good for them. The average A-Life researcher can offer biologists only three things that other biologists cannot: (1) programming skills and computer access, (2) knowledge of relevant A-Life work, and (3) a potentially useful naivete about biology that forces biologists to explain things clearly and to confront any hidden assumptions they may have.

From the biologists' perspective, A-Life work is too often a high-risk, low-gain enterprise. A-Life collaborators must show sufficient biological knowledge that the biologist feels confident she or he will not waste a year supporting a naive, over-complex, un-analyzable, un-comparable piece of hackwork.

3. Do a thorough scholarly review of the current biology relevant to the problem, especially formal models and conceptual issues

Collaborators can identify and summarize relevant biological literature, but there is no substitute for reading the key recent papers oneself. This step may seem like drudgery. But it has three consolations: theoretical biologists generally write clearly and even wittily, the mathematics is generally not too arcane, and every important paper read can save one from weeks of futile labor on already-solved issues.

There are three places to start: online library databases, good biology textbooks, and recent review monographs. Databases such as BIOSIS make good scholarship much easier. The full abstracts of every paper from recent issues of every major biological journal are available on the Internet. They can be accessed by subject, keywords, title, author, or institution. Many full papers can be accessed directly by ftp from the authors' home computer. Trudging to the library is necessary only after one knows exactly what one is looking for.

Compared to other fields, biology textbooks are unusually useful as resources, because recent, advanced work can be described simply enough for students to understand. Basic texts on evolution include Darwin (1859, 1871), Futuyama (1986), and Ridley (1993); see Sigmund (1993) for a more simulation-oriented introduction. The following list offers some of my other favorite resources in evolutionary biology, organized by topic: animal behaviour (Alcock, 1989), animal play behaviour (Fagen, 1981), behavioural ecology (Krebs & Davies, 1991), coevolution (Futuyama & Slatkin, 1983),

comparative method (Harvey & Pagel, 1991), cultural evolution (Boyd & Richerson, 1985; Cavalli-Sforza & Feldman, 1981), evolutionary game theory (Maynard Smith, 1982), evolutionary innovation (Nitecki, 1990), evolutionary psychology (Barkow et al., 1992; Buss, 1994; Daly & Wilson, 1988; Tooby & Cosmides, 1990), foraging theory (Stephens & Krebs, 1986), history of evolutionary theory (Cronin, 1991; Kohn, 1985; Richards, 1987), host-parasite coevolution (Ewald, 1993), human evolution (Foley, 1987; Mellars & Stringer, 1989), life history theory (Roff, 1992; Stearns, 1992), macroevolution (Eldredge, 1989), origins of sex (Williams, 1975; Maynard Smith, 1977; Michod & Levin, 1988), parental care (Clutton-Brock, 1991), protean behaviour (Driver & Humphries, 1988), sexual selection (Andersson, 1994; Daly & Wilson, 1983; Ridley, 1993), self-organization (Kauffman, 1993), social evolution (Byrne & Whiten, 1988; Dunbar, 1988; Trivers, 1985; Wilson, 1975), and theoretical population genetics (Gale, 1990).

4. Develop a well-targeted simulation that extends current biological models and yields directly comparable results

Historically, theoretical biology had to choose between verbal models that try to describe and explain organic complexity, and mathematical or spatial models that try to capture idealized but important aspects of that complexity. These two tools alone were sufficient to develop a quite satisfactory account of life and its evolution. But some of the most interesting phenomena and processes have always defied both verbal description and formal analysis.

For example, mathematical population genetics models of evolution must often make strong and unrealistic assumptions to yield tractable equations. Genetics are usually modelled by only a few loci of two alleles each, typically without linkage (correlations between alleles on different genes), dominance (of one allele by another given heterozygosity), or epistasis (nonlinear interactions between genes during development.) Populations are assumed to be of infinite size, with random mating,

uniform age and sex structure, and no geographic structure. Phenotypes are usually skipped entirely, without attending to life-history, learning, or contingent behaviour. The results of such models can be important in understanding simple evolutionary dynamics, and in exploring the implications of hypotheses about those dynamics. But such models are very weak at coping with phenomena such as complex phenotypes, flexible behaviour, co-evolution, or evolutionary innovation.

A powerful way of using A-Life simulations is to take an existing formal model from theoretical biology and relax the assumptions (preferably one at a time) that were required to make the mathematics tractable. The results of such a simulation are then directly comparable to the results of the existing formal model, and will be comprehensible and relevant for biologists. This strategy has been used to good effect by Collins and Jefferson (1992) to model runaway sexual selection, by De Bourcier and Wheeler (1994) to model aggressive signalling, by Webb (1994) to model cricket phonotaxis, and by Boekhorst and Hogeweg (1994) to model orangutan foraging.

However, theoretical biologists already use simulation in this way quite regularly. A-Life must go further, by exploiting the strengths of representing biological processes as computational procedures that allow evolution and emergence (see section 4), while ensuring that the results of such evolution and emergence remain comparable to those from current mathematical models in biology.

5. Explore cause and effect in the simulation by running systematic comparisons across different conditions

As computers become faster and easier to use, the temptation grows ever stronger to run tomorrow's computer simulation before yesterday's has been fully understood. Analysis is generally more tedious than synthesis. This seems particularly true in A-Life, which offers such unique prospects for synthetic demonstrations that are visually dazzling, emotionally engaging, and

February 1, 1995

11

intellectually exciting. It is also more fun to speculate about simulation's metaphysics than its methodology: philosophizing about A-Life has its own seductions and snares that distract attention from A-Life's real potentials and problems as science.

To ensure A-Life's relevance as theoretical biology, we must develop better methods for measurement and experimentation in our simulations. Whereas graphics may briefly catch the eyes of theoretical biologists, only solid experiments will win over their hearts. Our observational and experimental methods will have to approach the sophistication of taxonomy, comparative biology, ethology, psychology, ecology, and evolutionary biology. And we will have to develop new methods, because simulation can yield data that empirical biology cannot.

Early A-Life research consisted largely of proof-of-concept demonstrations: local interactions can lead to certain emergent effects suggestive of biological systems. Such results were important in establishing the concepts of self-organization, emergence, and bottom-up, decentralized control. Mature A-Life must put these concepts to scientific use through systematic experimentation. This means manipulating independent variables and observing effects on dependent variables, with both kinds of variables representing biologically meaningful parameters. A-Life researchers should first identify a research question, then identify comparisons and measurements of some system that could resolve the question, and only then should the simulation itself be designed. More often, researchers get interested in some biological phenomenon, spend months writing simulation code that aims to capture salient features of the phenomenon, and only later worry about what runs to do and what data to plot.

Evolutionary biology has been limited by its limited repertoire of data: fossils, short-term breeding experiments, and observations and experiments on living species (including their genotypes, phenotypes, and behaviours). From such sketchy, short-term, or static patterns, evolutionists try to

infer the dynamics of long-term processes. Such processes could not be studied directly in replicable, controlled experiments that could yield large amounts of relevant data.

Simulation allows access to much richer information as evolution progresses. Data falls into three rough categories. *Explicit simulation data* records any variables and data structures explicitly used by the simulation program; *Observed emergent data* records emergent effects using special measurement algorithms as the simulation progresses; *Experimentally derived data* requires on-line or post-hoc experiments to be performed on simulated things. For example, in an evolutionary simulation, genotypes would be explicit simulation data (since they are data structures in the simulation program). Phenotypes and behaviours would be observed emergent data (since they are constructed through generative development and emitted from a complex control system), and evolutionary dynamics would also be observed emergent data (since they arise from the unpredictable interactions of many agents over many generations). One could run post-hoc contests between early-generation individuals and later-generation individuals to yield experimentally derived data (see Cliff & Miller, in press). Each of these data types can be subject to statistical and comparative analysis on-line or post-hoc. Statistical summary data will become increasingly important because large simulations easily produce many megabytes of raw data: For example, one could easily record complete genotypes and phenotypes for every individual from every generation, complete records of each individual's lifetime interactions with the environment and other individuals, including all movement and behaviour, and all state variables such as sensory input, motor output, internal state, and energy state, and complete records of reproductive success, matings, genetic operator results, and lineages. A simulation that takes only a few hours to run may deserve weeks or months of analysis and replication.

Debugging simulation code may also require sophisticated measurement tools and experiments. With complex programs that produce emergent behaviour rather than checkable answers to

February 1, 1995

13

specific problems, it is not trivial to know whether the code is working appropriately. Making sure the code compiles and the program doesn't crash is just the first step. With humbling frequency, a simulation that seems to produce reasonable and interesting data under one set of conditions will produce, under slightly different conditions, weeks later (and often immediately before a conference), anomalous output that highlights some hidden, tiny, critical bug that invalidates weeks of results. There are three strategies for avoiding this.

First, develop and monitor a number of standard diagnostics that represent the low-level flow of control through the simulation code and the high-level flow of genes, information, matter, energy, and agents through the simulation world. Such measurements may be distinct from the data used to make scientific points about simulation results, and should be used across a wide array of simulations over a substantial period of time, so that diagnostics indicative of good code can be learned and anomalies can be registered quickly and easily. Second, run a set of standard control experiments specifically designed to test the robustness of the simulation code, not to test scientific hypotheses. Since most bugs happen at the interface between otherwise workable modules of code, the system as a whole should be assessed using a variety of simple tests. Third, write the code so that it can be put in the public domain and used by any interested party. This helps make the code and comments clear enough that they could be understood by anyone: things that are universally comprehensible rarely turn out to be bugs. But this also allows a given piece of code to be put to a variety of uses in a variety of contexts, such that checking its reliability becomes an interactive social process rather than an introspective black art of second-guessing where one might have made programming errors.

6. *Publish the results in real biology journals, subject to peer review by real biologists*

The new journals *Adaptive Behavior* and *Artificial Life*, exciting though they are as symbols of A-Life's arrival, should be viewed as cozy playgrounds in which ideas and methods grow stronger before venturing out into the harsh world of real science. If a piece of A-Life research is good biology, it should be published in a good peer-reviewed biology journal such as *American Naturalist*, *Animal Behavior*, *Behavioral Ecology and Sociobiology*, *Behavior*, *Ecology*, *Ethology and Sociobiology*, *Evolution*, *Genetica*, *Journal of Comparative Physiology*, *Journal of Theoretical Biology*, *Quarterly Review of Biology*, *Theoretical Population Biology*, or *Trends in Ecology and Evolution*. If a piece of A-Life research (e.g. a model of cultural evolution or communication), is good psychological or social science, it should be published in an appropriate journal such as *Behavioral and Brain Sciences*, *Current Anthropology*, or *Psychological Review*. A-Life work of high quality and more general interest could be submitted to a premier journal such as *Nature* or *Science*. Ideally, an A-Life project would go through several iterations and refinements over several years, with publication first in some conference proceedings, then in our specialist journals *Adaptive Behavior* or *Artificial Life*, then in an appropriate peer-reviewed mid-level biology journal, and finally in a top-level biology journal such as *Animal Behavior*, *Evolution*, or *Journal of Theoretical Biology*.

The burden of proof is on us to explain our results to biologists in their own language and in their own journals. Imagine you are a busy evolutionary theorist already inundated with exciting, new, well-written papers from old friends and known peers – papers that fit perfectly into the ongoing stream of research by citing relevant work, addressing current unsolved issues, developing new ideas, and making clear advances. Would you really bother seeking out work in our specialist conference proceedings and journals, which must seem suspiciously inbred havens for our pretentious, naive, and largely irrelevant simulations? If A-Life work cannot withstand peer review by biologists, why should

biologists bother to read it?

4 What distinguishes A-Life from other biological simulation?

Simulation is no stranger to biology. Theoretical biology papers that include simulations appear regularly in journals such as *Animal Behavior*, *Evolution*, *Journal of Theoretical Biology*, and *Nature*. Several journals are heavily biased towards biological simulation, such as *Computer Applications in the Biosciences*, *Computer Methods and Programs in Biomedicine*, *Ecological Modelling*, *Health Physics*, *Journal of Computer-Aided Molecular Design*, and *Mathematical Biosciences*. What does A-Life offer that theoretical biology does not already have?

Traditional mathematics and simulations in theoretical biology try to capture self-organizational or evolutionary dynamics directly in equations or simple procedures that aim straight for the collective, emergent level. Such methods are only tractable when they implicitly represent the components of biological systems as simple, stable, homogeneous, and predictable. Equations don't generally allow surprising, emergent behaviour – and neither do simulations based on equations.

The key advance in A-Life has been to allow 'emergence' by (1) representing individual biological things explicitly as computational procedures, (2) allowing self-organizational and evolutionary processes to emerge spontaneously from these things, and (3) making observations and measurements about the resulting patterns and dynamics at the individual and collective levels (see Taylor & Jefferson, 1994). This advance has allowed A-Life to explore the interaction of many biological units (molecules, cells, organisms, or populations) at several levels of description over different time-scales (behaviour, development, or evolution).

A-Life has used emergence in two main ways: as a proof-of-concept to show that certain biological phenomena can arise from distributed interactions among many local components, or as an extension to make current theoretical biology models more complete and realistic. The first way has led to some messianic predictions that theoretical biology will be revolutionized, perhaps with emergence replacing evolution as the central explanatory principle of life. The second way is more conservative and, I think, more useful: it embraces emergence without getting obsessed with it. Kauffman's (1993) research, despite its Emergentist rhetoric, seems a powerful example of this second strategy.

A-Life simulations have other advantages. Conceptually, the requirements of programming force researchers to make assumptions explicit and processes computable; formal equations are actually quite ambiguous compared to computer programs. Computationally, the speed of modern hardware allows biological systems to be simulated at levels of complexity unimaginable only a couple of decades ago. Experimentally, simulation offers a high degree of control, flexibility, and replicability. Analytically, simulation allows accurate measurement of very large amounts of data, and automatic statistical analysis and visualization of that data. Socially, simulation code can be shared over computer networks, promoting easy replication and extension of results.

There is a continuum between A-Life as high-level theoretical biology, studying general processes and patterns of evolution, and A-Life as empirical biology, modelling specific data from certain taxa. Some A-Life models the mechanisms or effects of a certain set of behaviours from a single species; but many biologists already develop such models routinely in their empirical work. More usefully, A-Life can model more general classes of behaviours shared across many species, such as courtship, cooperation, pursuit and evasion, communication, collective behaviour, or flocking. Models that predict different varieties of the behaviour for different species under different conditions would prove especially valuable to empirical biologists.

5 Open questions in theoretical biology that A-Life might usefully address

Given these strengths of A-Life, we can turn now to ask what open questions in theoretical biology might be especially amenable to A-Life simulation. The following areas seem promising: evolutionary innovations, interactions between different adaptive processes and different selective forces, origins and effects of mental and behavioural adaptations, and "life as it could be", logically and extra-terrestrially.

Theoretical biology cannot yet explain major evolutionary innovations such as the evolution of life (Eigen & Schuster, 1985), sex (Williams, 1975; Margulis & Sagan, 1986; Maynard Smith, 1978; Michod & Levin, 1988), multi-cellular bodies (Buss, 1987), or the human brain (Miller, 1993; Ridley, 1993). Explaining even medium and minor innovations has proven difficult (Nitecki, 1990). Some innovations probably result from interplays between phenotypic self-organization and natural selection (Kauffman, 1993). The powerful, capricious process of sexual selection may also promote evolutionary innovations (Eberhard, 1985; Miller, 1994; Miller & Todd, in press; Ryan, 1990). A-Life research is uniquely appropriate for discovering further principles about how various forms of self-organization and selection interact in biologically realistic systems so as to produce evolutionary innovations. Before any further simulations are attempted in this area, it would be worth reading the classic work on evolutionary innovations by Romanes (1897), Weismann (1917), Simpson (1953), Mayr (1960), and Nitecki (1990).

A-Life is also suited to studying interactions between different adaptive processes such as behaviour, learning, development, and evolution, which operate on different time-scales according to different dynamics (see Eldredge, 1985). Many theorists have suggested that development and evolution interact in important ways (Bonner, 1982; Gould, 1977; Raff & Raff, 1987), but it has proven difficult to go beyond verbal hypotheses. Some A-Life work focuses on such interactions (Dawkins,

February 1, 1995

18

1989; Kauffman, 1993). Explicit simulation of the evolution of generative development and its effects on the further course of evolution would be valuable. Likewise, Baldwin (1896) suggested that learning facilitates evolution, and some simulations have confirmed this 'Baldwin effect' (Hinton & Nowlan, 1985; Ackley & Littman, 1991; French & Messinger, 1994; see also Todd & Miller, 1993), but much more research is needed. A-Life may also continue to illuminate the interactions between genetic evolution and 'cultural evolution' (Boyd & Richerson, 1985; Cavalli-Sforza & Feldman, 1981), but anyone tempted by this topic is strongly advised to read Tooby and Cosmides (1992) and Hirschfeld and Gelman (1994) as an antidote against over-estimating the role that 'culture' plays in shaping human behaviour and cognition.

Theoretical biology has also found it difficult to understand how selective forces interact under realistic conditions. In many cases, there are complex interactions between natural selection and sexual selection, or between habitat choice, ecology specialization, social behaviour, and sexual behaviour. Formal models must generally assume that one of these factors dominates (e.g. given an econiche, what social behaviour follows?), but in most cases it is not clear what causes what: do lineages adapt their social behaviour to their habitat, or do they choose habitats appropriate for their social behaviour? Ultimately, biology needs a theory about the relative speed at which different kinds of adaptations evolve, because the slower-evolving kinds will set the context for the faster-evolving kinds. It seems reasonable to assume that behaviours usually evolve faster than bodies, and body forms evolve faster than basic life-history parameters, but A-Life simulations might tell us why.

The evolution of cognition and behaviour is extremely difficult to study with formal models, because the phenotypes are so complicated and have such unexpected evolutionary effects. We know almost nothing about the detailed evolution of systems for sensation, perception, cognition, learning, memory, decision-making, motivation, emotion, cyclical behaviour, and motor control. Neuroethology

and comparative psychology only reveal the outcomes of neural and cognitive evolution, and fossils do not provide details of neural circuits during evolution. A-Life systems that explicitly simulate the evolution of nervous systems interacting with each other and with complex environments may be our only hope for constructing theories of mental and behavioural evolution.

Once behavioural and cognitive adaptations have evolved, they can influence the further course of evolution within and across species. The role of 'mind' in guiding evolution has usually been overlooked entirely, or conflated with a mystical, progressivist, animism as in the work of Herbert Spencer (1855; see Godfrey-Smith, 1994) and William McDougall (1929; see Boden, 1972). Only a few biologists developed a Darwinian view of minds as selective forces (Morgan, 1888; Thompson, 1920; see Richards, 1987). But very recently, several theoretical biologists have begun to recognize the importance of perception and cognition as selective forces in the evolution of diverse phenomena such as camouflage, mimicry, warning colouration, sexual ornaments, flowers, fruits, reciprocal altruism, kin selection, deception, self-deception, social intelligence, protean behaviour, signals, communication, and language (Barth, 1991; Bateson, 1988; Byrne & Whiten, 1988; Dawkins & Krebs, 1978; Driver & Humphries, 1988; Endler, 1992; Guilford & Dawkins, 1991; Miller, 1993; Miller & Cliff, 1994; Miller & Freyd, 1993; Plotkin, 1988; Ryan, 1990). This area is uniquely appropriate for A-Life simulations in which sensory-motor capacities are both subjects and objects of selection. Sexual selection through mate choice is a particularly clear case of such co-evolution: interactions between mate choice mechanisms and sexually-selected traits are poorly understood, but have produced a stunning diversity of extravagant ornaments, behaviours, and innovations (Eberhard, 1985; Miller, 1993, 1994; Ridley, 1993). Only a few A-Life researchers have explored sexual selection so far (Collins & Jefferson, 1992; Miller & Todd, 1993; Todd & Miller, 1991, 1993). The dynamics of sexual selection are probably similar to those of other examples of 'psychological selection' (Miller, 1993). Biology is wide open for simulations that allow co-evolution between perceptual systems and

February 1, 1995

20

the things that they perceive.

A-Life has long paid lip service to the idea of putting terrestrial biology in a broader context by exploring "life as it could be" in the computer. But there has been very little real research in the two most obvious areas for such work: exploring the logical implications of alternative biologies, and exploring the likely biologies of extra-terrestrial environments.

The former simply requires seeing what happens if some apparently fixed parameter is set at an unconventional value. What if there were four sexes? What if Lamarckian inheritance were possible (see Ackley & Littman, 1994)? What if carnivorous plants could evolve psycho-kinesis? What if an animal's gonads or sensors could move independently of the rest of its body? What if an animal's "extended phenotype" (Dawkins, 1982) could reach into the past or the future? The sillier such questions seem, the more usefully they might undermine our unconscious assumptions about biology, and the more they might stretch our imaginations. A-Life, despite its pretenses to expanding the scope of possible biologies, has been tied far too tightly to real biology.

The latter area I would call 'computational xenobiology' if I didn't loathe neologisms. It requires setting up some simulation environment based on what is known about some extra-terrestrial habitat, and seeing what might evolve there given those physical constraints. For example, what kind of bodies might evolve inside a gas giant, in a star's corona, or on a neutron star? The basic principle of natural selection applies to any system of replicators anywhere in the universe (Dawkins, 1976), so A-Life might give us insights into the frequency, distribution, and variety of extra-terrestrial life. It would be comforting for A-Life to leave us not entirely surprised should Contact occur.

6 Does Strong A-Life allow stronger theoretical biology research than Weak A-Life?

The debate over ‘strong A-Life’ (computer processes as realizations of living systems) versus ‘weak A-Life’ (computer processes as simulations of living systems) can be extended in a methodological direction by asking: would it make any difference to theoretical biology if an A-Life system were construed as a realization rather than a simulation? Clearly, empirical biology would be affected: we would have to add a sixth kingdom of life to the current five (see Margulis & Schwartz, 1988), and I suppose that databases of biological phylogenies would have to be updated every time a new Ph.D. thesis in A-Life was written. Also granted is that acceptance of strong A-Life would imply that ‘life’ (like ‘mind’) is a functional, emergent property of certain systems with certain internal relations, external relations, and evolutionary histories (see Millikan, 1984). But my focus is on biological theory: what could we learn about life and evolution from *doing* strong A-Life that we could not learn from doing weak A-Life?

In answering this question, we must appreciate that many theoretical biologists view instances of real terrestrial life as little more than the outcomes of ‘simulations by Nature’ to inform them about how evolution works. That is, empirical biology constrains biological theory in almost the same way that simulation does, so for some theoretical biologists, even real terrestrial biology could be considered ‘weak A-Life’. Life-forms that result from artificial selection by human breeders or genetic engineering also blur the distinction between realization and simulation. If one makes a strong division between Nature and Culture, such life-forms are experimental simulations of what *would* happen if a lineage were subjected to some selective pressure or mutation in Nature; from a more integrated perspective, such life-forms are simply the outcome of a thoroughly Natural process that happens to include humans as selective forces. Likewise, experimental biology research that records animal behaviour in unnatural laboratory conditions could be viewed either as ‘realizations of behaviour *in*

February 1, 1995

22

extremis', or 'simulations of natural behaviour'. Does it make any difference to theoretical biology either way?

Consider an issue in evolutionary theory that might be solved in two ways. Ms. Goodmaths develops a mathematical population genetics model that represents changes in gene frequencies using differential equations, whereas Mr. Badmaths programs an A-Life model that represents genes themselves in a genetic algorithm. In the latter case, one might argue that the genes are 'alive' in the limited sense of replicating because they really are copied within computer data structures. Suppose the two models are both good and give the same answer. The 'aliveness' of the genes in the A-Life model is simply irrelevant to the theoretical biology. If the models are formally equivalent, one can move smoothly from the differential equations through discrete iterative approximations to genetic algorithms without affecting the results in the slightest. The evolutionary dynamics have 'multiple realizability' in formal equations, computational approximations, and procedural simulations. (Only if Mr. Badmaths' simulation goes beyond the complexity that Ms. Goodmaths' equations can represent, is there any point in doing the simulation.)

An analogy to high-energy physics may also help here. The events that occur within colliders at CERN have an ambiguous status. Empirical physicists treat them as 'real physical events' that reveal forces operating outside colliders. But for theoretical physicists, colliders may as well be viewed as very special, very expensive computers that simulate the physics of the very early universe, shortly after the Big Bang. It does not really matter for theoretical physics whether collider events are viewed as 'realizations of current physics' or 'simulations of early-universe physics', because in both cases the problem remains of how to generalize to processes outside the collider.

For theoretical biologists who accept a functionalist view of life, data from ‘natural’ biology, artificial selection, genetic engineering, weak A-Life, and strong A-Life should have the same status: they all reveal things about biological processes in one system (e.g. species or simulation) that may or may not generalize to other systems. We need better methodologies to guide such generalizations, based partially on results from theoretical biology itself. Darwin’s theory of descent with modification guides generalization across lineages, and his theory of adaptation through natural selection guides generalization across habitats and niches. Ecological learning theory (Davey, 1989; Gallistel, 1990; see Miller & Todd, 1990) suggests that it is dangerous to generalize about ‘learning’ across different tasks and species. New theories in evolutionary psychology are showing that generalizations across cognitive domains are often perilous (Hirschfeld & Gelman, 1994; Tooby & Cosmides, 1992). In their role as philosophers of biology, philosophers of A-Life might help clarify how to generalize about biological processes inside and outside computers.

A great deal turns on the Strong A-Life debate: public interest in A-Life, research funding for A-Life, and the social status of A-Life pioneers (God-like disseminators of novel life-forms, or sad hackers in need of a real life?). But for theoretical biology research, I believe the debate is simply an irrelevant distraction.

7 Conclusions

If Darwin were alive today, he would probably be enthusiastic about A-Life research, especially about open-ended simulations of evolution. Simulation offers evolutionary biologists, behavioural ecologists, evolutionary psychologists, and ethologists an unprecedented opportunity for understanding complex evolutionary dynamics that defy traditional analytic methods such as those used in mathematical population genetics (Gale, 1990), evolutionary game theory (Maynard Smith, 1982), and optimization theory (Mangel & Clark, 1988; Schoemaker, 1991; see Miller, 1991; Miller &

February 1, 1995

24

Todd, 1991). Such simulations will probably become the most important scientific tool for understanding evolution since Darwin (1859, 1871) first developed the theories of natural and sexual selection.

But A-Life research will only be as good as A-Life research methods. Synthesis without analysis is engineering rather than science (and most current engineering methods include much analysis). Observation without experimentation is just cyber-tourism and cheap thrills. To increase the relevance and utility of A-Life to theoretical biology, six heuristics were suggested: (1) pick a known, unsolved problem in theoretical biology, (2) collaborate with real biologists, (3) do thorough biological scholarship, (4) develop a simulation that extends current biological models, (5) use simulation experiments to explore cause and effect, and (6) publish in real peer-reviewed biology journals. I then identified some areas of theoretical biology that contain a large number of open questions that could be addressed with A-Life simulations. These areas include the origin of evolutionary innovations through self-organization and selection; the interaction of different adaptive processes (behaviour, learning, development, evolution); the evolutionary origins and effects of behaviour, perception, and cognition; the interaction of multiple selective forces; and "life as it could be", theoretically and extra-terrestrially. Finally, I argued that the debate over 'strong A-Life' is largely irrelevant to A-Life's uses in theoretical biology.

The prospects for A-Life as real biological science are exciting enough; the field need not wait for philosophers to give their certificate of 'Real Life' authenticity to A-Life. On the contrary, I believe the best philosophy of A-Life will be a modest, methodological philosophy of science, rather than a metaphysical philosophy of life.

Acknowledgments

This research was supported by NSF-NATO Post-Doctoral Research Fellowship RCD-9255323 and NSF Research Grant INT-9203229. Thanks to Maggie Boden, Dave Cliff, Inman Harvey, Peter Todd, and Michael Wheeler for useful discussions.

References

- Ackley, D., & Littman, M. (1992). Interactions between learning and evolution. In C. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II*, pp. 487-509. New York: Addison-Wesley.
- Ackley, D., & Littman, M. (1994). A case for Lamarckian evolution. In C. Langton (Ed.), *Artificial Life III*, pp. 3-10. New York: Addison-Wesley.
- Alcock, J. (1993). *Animal behavior: An evolutionary approach* (5th Ed.). Sunderland, MA: Sinauer Associates.
- Andersson, M. B. (1994). *Sexual selection*. Princeton U. Press.
- Barkow, J. H., Cosmides, L., & Tooby, J. (Eds.). (1992). *The adapted mind: Evolutionary psychology and the generation of culture*. Oxford U. Press.
- Barth, F. G. (1991). *Insects and flowers: The biology of a partnership*. Princeton U. Press.
- Bateson, P. (1988). The active role of behavior in evolution. In M.-W. Ho & S. W. Fox (Eds.), *Evolutionary processes and metaphors*, pp. 191-207. New York: John Wiley.
- Bedau, M., & Packard, N. (1992). Measurement of evolutionary activity, teleology, and life. In C. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II*, pp. 431-461. New York: Addison-Wesley.
- Boden, M. (1972). *Purposive explanation in psychology*. Harvard U. Press.

- Boekhorst, I. J. A. te, & Hogeweg, P. (1994). Effects of tree size on travelband formation in Oran-
Utans: Data analysis suggested by a model study. In R. A. Brooks & P. Maes (Ed.), *Artificial
Life IV*, pp. 119-129. MIT Press/Bradford Books.
- Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. U. Chicago Press.
- Buss, D. M. (1994). *The evolution of desire: Human mating strategies*. New York: Basic Books.
- Buss, L. W. (1987). *The evolution of individuality*. Princeton U. Press.
- Byrne, R., & Whiten, A. (Eds.). (1988). *Machiavellian intelligence: Social expertise and the evolution
of intellect in monkeys, apes, and humans*. Oxford U. Press.
- Cariani, P. (1992). Emergence and artificial life. In C. Langton, C. Taylor, J. D. Farmer, & S. Ras-
mussen (Eds.), *Artificial Life II*, pp. 775-797. New York: Addison-Wesley.
- Cavalli-Sforza, L. L., & Feldman, M. W. (1981). *Cultural transmission and evolution*. Princeton U.
Press.
- Cliff, D. (1991). Computational neuroethology: A provisional manifesto. In J. A. Meyer & S. W. Wil-
son (Eds.), *From animals to animats: Proceedings of the First International Conference on
Simulation of Adaptive Behavior*, pp. 29-39. MIT Press.
- Cliff, D., & Miller, G. F. (in press). Co-Evolution of pursuit and evasion II: Simulation methods and
results. For *Adaptive Behavior*.
- Clutton-Brock, T. H. (1991). *The evolution of parental care*. Princeton U. Press.
- Collins, R. J., & Jefferson, D. R. (1992). The evolution of sexual selection and female choice. In F. J.
Varela and P. Bourguine (Eds.), *Toward a practice of autonomous systems: Proceedings of the
First European Conference on Artificial Life*, pp. 327-336. MIT Press.
- Cronin, H. (1991). *The ant and the peacock: Altruism and sexual selection from Darwin to today*.
Cambridge U. Press.
- Daly, M., & Wilson, M. (1983). *Sex, evolution, and behavior* (2nd Ed.). Boston: Willard Grant Press.

- Daly, M., & Wilson, M. (1988). *Homicide*. New York: Aldine.
- Darwin, C. (1859). *On the origin of species by means of natural selection*. London: John Murray.
- Darwin, C. (1862). *On the various contrivances by which orchids are fertilized by insects*. London: John Murray.
- Darwin, C. (1871). *The descent of man, and selection in relation to sex* (2 Vols.). London: John Murray.
- Davey, G. (1989). *Ecological learning theory*. London: Routledge.
- Dawkins, R. (1976). *The selfish gene*. Oxford U. Press.
- Dawkins, R. (1982). *The extended phenotype*. Oxford: W. H. Freeman.
- Dawkins, R. (1989). The evolution of evolvability. In Langton, C. (Ed.), *Artificial Life*, pp. 201-220. New York: Addison-Wesley.
- Dawkins, R., & Krebs, J. R. (1978). Animal signals: information or manipulation? In J. R. Krebs & N. B. Davies (Eds.), *Behavioral ecology: An evolutionary approach* (2nd Ed.), pp. 282-309. Oxford: Blackwell Scientific.
- De Bourcier, P., & Wheeler, M. (1994). Signalling and territorial aggression: An investigation by means of synthetic behavioural ecology. In D. Cliff, P. Husbands, J. A. Meyer, & S. Wilson, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pp. 463-472. MIT Press.
- Driver, P. M., and Humphries, N. (1988). *Protean behavior: The biology of unpredictability*. Oxford U. Press.
- Dunbar, M. (1988). *Primate social systems*. London: Croom Helm.
- Eberhard, W. G. (1985). *Sexual selection and animal genitalia*. Harvard U. Press.
- Eigen, M., & Schuster, P. (1985). Stages of emerging life – five principles of early organization. *J. Molecular Evolution*, 19, 47-61.

- Eldredge, N. (1985). *Unfinished synthesis: Biological hierarchies and modern evolutionary thought*. Oxford U. Press
- Eldredge, N. (1989). *Macroevolutionary dynamics: Species, niches, and adaptive peaks*. New York: McGraw-Hill.
- Endler, J. A. (1992). Signals, signal conditions, and the direction of evolution. *American Naturalist*, 139, S125-S153.
- Ewald, P. (1993). *The evolution of infectious disease*. Oxford U. Press.
- Fagen, R. (1981). *Animal play behavior*. Oxford U. Press.
- Fisher, R. A. (1930). *The genetical theory of natural selection*. Clarendon Press.
- French, R. M., & Messinger, A. (1994). Genes, phenes, and the Baldwin Effect: Learning and evolution in a simulated population. In R. A. Brooks & P. Maes (Ed.), *Artificial Life IV*, pp. 277-282. MIT Press/Bradford Books.
- Foley, R. (1987). *Another unique species: Patterns in human evolutionary ecology*. Harlow, Essex, England: Longman Scientific & Technical.
- Futuyama, D. J. (1986). *Evolutionary biology* (2nd Ed.). Sunderland, MA: Sinauer.
- Futuyama, D. J. & Slatkin, M. (Eds.). (1983). *Coevolution*. Sunderland, MA: Sinauer.
- Gale, J. S. (1990). *Theoretical population genetics*. London: Unwin Hyman.
- Gallistel, C. R. (1990). *The organization of learning*. MIT Press.
- Godfrey-Smith, P. (1994). Spencer and Dewey on life and mind. In R. A. Brooks & P. Maes (Ed.), *Artificial Life IV*, pp. 80-89. MIT Press/Bradford Books.
- Guilford, T., & Dawkins, M. S. (1991). Receiver psychology and the evolution of animal signals. *Animal Behavior*, 42, 1-14.
- Harvey, P. H., & Pagel, M. D. (1991). *The comparative method in evolutionary biology*. Oxford U. Press.

- Hirschfeld, L. A., & Gelman, S. A. (Eds.) (1994). *Mapping the mind: Domain specificity in cognition and culture*. Cambridge U. Press.
- Kauffman, S. (1993). *The origins of order: Self-organization and selection in evolution*. Oxford U. Press.
- Kohn, D. (Ed.). (1985). *The Darwinian Heritage*. Princeton U. Press.
- Krebs, J. R., & N. B. Davies (Eds.) (1993). *Behavioral ecology: An evolutionary approach* (3rd Ed.). Oxford: Blackwell Scientific.
- Langton, C. (1989). Artificial life. In Langton, C. (Ed.), *Artificial Life*, pp. 1-47. New York: Addison-Wesley.
- Lewin, R. (1993). *Human evolution* (3rd Ed.). Boston: Blackwell Scientific.
- Mangel, M., & Clark, C. W. (1988). *Dynamic modelling in behavioral ecology*. Princeton U. Press.
- Margulis, L., & Sagan, D. (1986). *Origins of sex: Three billion years of genetic recombination*. Yale U. Press.
- Margulis, L., & Schwartz, K. V. (1988). *Five kingdoms: An illustrated guide to the phyla of life on earth* (2nd Ed.). New York: W. H. Freeman.
- Maynard Smith, J. (1978). *The evolution of sex*. Cambridge U. Press.
- Maynard Smith, J. (1982). *Evolution and the theory of games*. Cambridge U. Press.
- Mayr, E. (1960). The emergence of evolutionary novelties. In S. Tax (Ed.), *Evolution after Darwin* (Vol. I), pp. 349-380. U. Chicago Press.
- McDougall, W. (1929). *Modern materialism and emergent evolution*. London: Methuen.
- Mellars, P., & Stringer, C. (Eds.). (1989). *The human revolution: Behavioral and biological perspectives on the origins of modern humans*. Edinburgh U. Press.
- Michod, R. E., & Levin, B. R. (Eds.). (1988). *The evolution of sex: An examination of current ideas*. Sunderland, MA: Sinauer.

- Miller, G. F. (1991). Two dynamic criteria for validating claims of optimality. *Behavioral and Brain Sciences*, 14(2): 228-229.
- Miller, G. F. (1993). *Evolution of the human brain through runaway sexual selection: The mind as a protean courtship device*. Ph.D. thesis, Stanford University Psychology Department. (To be published as a book by MIT Press/Bradford Books in 1995.)
- Miller, G. F. (1994). Exploiting mate choice in evolutionary computation: Sexual selection as a process of search, optimization, and diversification. In T. C. Fogarty (Ed.), *Evolutionary Computing: Proceedings of the 1994 Artificial Intelligence and Simulation of Behavior (AISB) Society Workshop*, pp. 65-79. Springer-Verlag.
- Miller, G. F., & Cliff, D. (1994). Protean behavior in dynamic games: Arguments for the co-evolution of pursuit-evasion tactics in simulated robots. In D. Cliff, P. Husbands, J. A. Meyer, & S. Wilson, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pp. 411-420 MIT Press.
- Miller, G. F., & Freyd, J. J. (1993). *Dynamic mental representations of animate motion: The interplay among evolutionary, cognitive, and behavioral dynamics*. Cognitive Science Research Paper CSRP-290, University of Sussex.
- Miller, G. F., & Todd, P. M. (1990). Exploring adaptive agency I: Theory and methods for simulating the evolution of learning. In D. S. Touretsky, J. L. Elman, T. J. Sejnowski, & G. E. Hinton (Eds.), *Proceedings of the 1990 Connectionist Models Summer School*, pp. 65-80. San Mateo, CA: Morgan Kaufmann.
- Miller, G. F., & Todd, P. M. (1991). Let evolution take care of its own. *Behavioral and Brain Sciences*, 14(1): 101-102.
- Miller, G. F., & Todd, P. M. (1993). Evolutionary wanderlust: Sexual selection with directional mate preferences. In J.-A. Meyer, H. L. Roitblat, & S. W. Wilson (Eds.), *From Animals to Animats*

- 2: *Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, pp. 21-30. MIT Press.
- Miller, G. F., & Todd, P. M. (accepted). The role of mate choice in biocomputation: Sexual selection as a process of search, optimization, and diversification. For W. Banzaf & F. Eeckman (Eds.), *Proceedings of the 1992 Monterey Biocomputation Workshop*.
- Millikan, R. (1984). *Language, thought, and other biological categories*. MIT Press.
- Morgan, C. L. (1888). Natural selection and elimination. *Nature*, Aug. 16, 370.
- Nitecki, M. H. (Ed.). (1990). *Evolutionary innovations*. U. Chicago Press.
- Pattee, H. H. (1989). Simulations, realizations, and theories of life. In Langton, C. (Ed.), *Artificial Life*, pp. 63-77 New York: Addison-Wesley.
- Plotkin, C. G. (Ed.). (1988). *The role of behavior in evolution*. MIT Press.
- Raff, R. A., & Raff, E. C. (Eds.). (1987). *Development as an evolutionary process*. New York: Alan R Liss.
- Richards, R. J. (1987). *Darwin and the emergence of evolutionary theory of mind and behavior*. U. Chicago Press.
- Ridley, Mark (1993). *Evolution*. Boston: Blackwell Scientific.
- Ridley, Matt (1993). *The red queen: Sex and the evolution of human nature*. London: Viking.
- Roff, D. A. (1992). *The evolution of life histories: Theory and analysis*. New York: Chapman & Hall.
- Romanes, G. J. (1897). *Darwin, and after Darwin. II. Post-Darwinian Questions. Heredity and Utility* (2nd Ed.). Chicago: Open Court Publishing.
- Ryan, M. J. (1990). Sexual selection, sensory systems, and sensory exploitation. *Oxford Surveys of Evolutionary Biology*, 7, 156-195.
- Said, E. (1978). *Orientalism*. New York: Pantheon Books.
- Schoemaker, P. J. H. (1991). The quest for optimality: A positive heuristic for science. *Behavioral*

and Brain Sciences, 14(2): 205-215.

Sigmund, K. (1993). *Games of life: Explorations in ecology, evolution, and behavior*. Oxford U. Press.

Simpson, G. (1953). *The major features of evolution*. Columbia U. Press.

Sober, E. (1992). Learning from functionalism – Prospects for strong artificial life. In C. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II*, pp. 749-765. New York: Addison-Wesley.

Spencer, H. (1855). *Principles of psychology*. London: Longman, Brown, & Green.

Stearns, S. C. (1992). *The evolution of life histories*. Oxford U. Press.

Stephens, D. W., & Krebs, J. R. (1986). *Foraging theory*. Princeton U. Press.

Taylor, C., & Jefferson, D. (1994). Artificial Life as a tool for biological inquiry. *Artificial Life*, 1/2, 1-13.

Thompson, A. (1920). *The system of animate nature* (2 Vols.). London: Williams & Norgate.

Todd, P. M., & Miller, G. F. (1991). On the sympatric origin of species: Mercurial mating in the Quicksilver Model. In R. K. Belew & L. B. Booker (Eds.), *Proceedings of the Fourth International Conference on Genetic Algorithms*, pp. 547-554. San Mateo, CA: Morgan Kaufmann.

Todd, P. M., & Miller, G. F. (1993). Parental guidance suggested: How parental imprinting evolves through sexual selection as an adaptive learning mechanism. *Adaptive Behavior*, 2(1): 5-47.

Tooby, J., & Cosmides, L. (1990). The past explains the present: Emotional adaptations and the structure of ancestral environments. *Ethology and sociobiology*, 11(4/5), 375-424.

Tooby, J., & Cosmides, L. (1992). The psychological foundations of culture. In J. H. Barkow, L. Cosmides, & J. Tooby, J. (Eds.), *The adapted mind: Evolutionary psychology and the generation of culture*, pp. 19-136. Oxford: Oxford U. Press.

Trivers, R. (1985). *Social evolution*. Menlo Park, California: Benjamin/Cummings.

- Webb, B. (1994). Robotic experiments in cricket phonotaxis. In D. Cliff, P. Husbands, J. A. Meyer, & S. Wilson, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pp. 45-54. MIT Press.
- Weismann, A. (1917). The selection theory. In *Evolution in modern thought, by Haeckel, Thomson, and Weismann, and Others*, pp. 23-86. New York: Boni and Liveright.
- Williams, G. C. (1966). *Adaptation and natural selection*. Princeton U. Press.
- Williams, G. C. (1975). *Sex and evolution*. Princeton U. Press.
- Wilson, E. O. (1975). *Sociobiology: The new synthesis*. Harvard U. Press.
- Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding, and selection in evolution. *Proc. Sixth Int. Congr. Genetics, 1*, 356-366.