

Presented at a Symposium on “The making up of organisms: mapping the future of biological models and theories,” Ecole Normale Supérieure, Paris, 9 June 2006.

The Evolution of Complexity

Mark A. Bedau

Reed College, Portland, Oregon, USA
ProtoLife SRL, Venice, Italy
European Center for Living Technology, Venice, Italy

The evolution of the biosphere exhibits a trend of increasing complexity of the most complex organisms. Even though we are uncertain about the proper way to measure complexity, it is hard to deny that the earliest prokaryotic cells are simpler than the eukaryotic cells that arose from them, and these were simpler than the multicellular life forms that evolved from them, and so on. But it is controversial how to interpret and explain this trend, and even how to describe it properly (see, e.g., Gould 1989, McShea 1991, Gould 1996, Sterelny and Griffiths 1999). Some think that increasing maximal organism complexity has for all intents and purposes been explained by Darwin, Miller and Urey, Watson and Crick, or some other past or present biologists. This chapter defends the alternative view that the trend is *not* yet adequately explained but instead is one of the major remaining challenges in understanding biological phenomena.¹

My argument will crucially employ a new explanatory scheme in the life sciences: what I will call “constructive” or “emergent” computer models, which are characteristic of “soft” artificial life. These models are useful in part because they can provide clear and decisive refutation of various hypotheses about how to explain the trend of increasing organism complexity. Even more important, in general only these models can provide clear and compelling explanations for the characteristic behavior of complex adaptive systems like the evolving biosphere. Thus, these models turn out to be exactly the right kind of tool for exploring how to explain the trend of increasing complexity; furthermore, they do so in a way that is readily subject to extension and empirical corroboration. However, these models are no panacea. They can be misleading or misinterpreted, and they have limitations. As we will see when we examine specific examples, the proper use of constructive models and the proper interpretation of their behavior requires care and experience, and creating especially insightful models is extremely challenging.

The arrow of complexity hypothesis

¹ Stuart Kauffman (2000) has done as much as anyone to call attention to this question. See also Farmer and Belin (1992) and Holland (1995).

It is useful to distinguish three things: (i) a *trend*, which is simply an observed directional change in some variable in some evolving system; (ii) a *robust regularity*, which is a generic or non-accidental trend (perhaps with exceptions), a statistical “law” about the time dynamics of some variable; and (iii) a *mechanism or process* that explains a trend, whether accidental or robust. So far we have simply been observing that there is a trend of increasing complexity of the most complex organisms. This observation does not necessarily imply that the trend is a robust regularity; it might simply be an accident. Nor is any particular mechanism or process implicated as the explanation of that trend. Our starting point is merely the observation of a trend.

But perhaps the trend is not merely an accident. Perhaps it is an instance of some general regularity. The *arrow of complexity hypothesis* is the hypothesis that evolution inherently creates increasingly complex adaptive organisms. This hypothesis is about the increasing complexity of the *most complex* organisms, not of all organisms or of life’s mean complexity. (A closely related hypothesis pertains to higher levels of organizations like ecosystems, but that is another topic.) Furthermore, the hypothesis concerns not mere complexity but *adaptive* complexity, that is, complexity that is an adaptation and serves some function. I take an adaptation to be a trait that comes to exist and persists in a lineage because it is beneficial for the organisms in the lineage, where this can be cashed out by saying that the trait is produced by natural selection for that trait (Bedau 1991, 1992). This view is in the spirit of the familiar etiological or historical approach to understanding biological functions (see, e.g., Allen et al. 1997).

The trend of increasing maximal organism complexity illustrates a clear and dramatic form of the emergence of biological complexity over time. Emergence had a checkered history in twentieth century philosophy because it was often construed as conflicting with forms of reductionism that were taken to be characteristic of naturalism. But notions of emergence that are compatible with naturalism and at least certain uncontroversial forms of reductionism have now been developed (e.g., Wimsatt 1997, Bedau 1997, 2003b, Boogerd et al. 2005). One would presumably want some such notion to characterize the emergence of increasing organism complexity because that trend is presumably a fully naturalistic process. The arrow of complexity hypothesis is the hypothesis that this dynamical emergent process is a generic property of some class of evolving systems, rather than just an accident.

The notion of complexity is notoriously difficult to define.² I will not propose any substantive definition here. When I talk about increasing complexity, I mean *qualitative* complexity; an organism with 100 vertebrae seems quantitatively more

² For a nice review of some of the conceptual problems, see Sterelny and Griffiths (1999).

complex than one with 10 vertebrae, but it does not seem qualitatively more complex. Something's qualitative complexity seems to be something roughly like the number of different kinds of fundamental properties it possesses. A vast number of different definitions of complexity have been proposed (e.g., Edmunds 1996, Lloyd 2001, Wolfram 2002). Some of the variety reflects different attempts to capture essentially similar intuitions about complexity, but many definitions attempt to capture fundamentally different intuitions about what complexity is. I will not attempt to navigate this conceptual thicket here, but will instead for simplicity rely on a vague but not wholly unintuitive notion of qualitative complexity of organisms. Since there are different notions of complexity, there are also correspondingly different hypotheses about the arrow of complexity. Similarly, there are a variety of analogous arrow hypotheses that concern not complexity but trends in increasing disparity, diversity, body mass, energy flow, among other notions. I will not here address whether the argument in this chapter generalizes to these other arrow hypotheses.

The arrow hypothesis should be understood as a robust regularity that has exceptions. We will see examples of evolving systems for which the arrow hypothesis fails altogether. Further, those systems for which the hypothesis holds are usually not always increasing in maximal organismal complexity; rather, sometimes their maximal complexity temporarily falls, for one reason or another. The hypothesis is that maximal complexity generically tends continually to rise. You could put it this way: For any time t with measured complexity c , there is some later time $t^* > t$ with complexity $c^* > c$. Any actual evolving system that we measure (such as the biosphere) at some finite time will by necessity be consistent with both the truth and falsity of the arrow of complexity hypothesis.

In addition, the hypothesis holds for evolving systems only under certain constraints, such as a continual supply of energy or material resources. There is no implication that maximal complexity must increase if those constraints fail. But the implication of the hypothesis is that the biosphere's increasing maximal complexity is not just an accident. The hypothesis does not specify the mechanism behind this (purported) robust regularity. It is rather analogous to the second law of thermodynamics, which simply states that there is an arrow of increasing entropy but does not specify the mechanism behind it. The key question about the arrow of complexity hypothesis is how to test whether it is true and, if it is, how to discover the mechanism behind it.

Replaying the tape of life

This issue of the creativity of evolution has been addressed by a number of biologists and philosophers (though not in exactly my terms). In *Wonderful Life* (1989), S. J. Gould proposed a famous procedure for testing things like the arrow of

complexity hypothesis. He called the test “replaying the tape of life.” To conceive of this test, one imagines a tape of the history of life, and one imagines rewinding and erasing that tape to some point in the early evolution of life. Then one imagines replaying the tape forward again, but this time having the course of life affected by different historical contingencies. In this way, what is accidental in the history of life might turn out differently, while what (if anything) is essential will remain the same. If one repeats this rewind-and-replay operation again and again and compares the results, robust regularities will stand out as statistical laws. Replaying the tape of life would be an excellent way to settle whether the arrow of complexity hypothesis is true.

Goald suggests that replaying the life’s tape would disprove the arrow of complexity hypothesis. He emphasizes that the evolution of life is a contingent historical process and is thus not law-like. He says that “almost every interesting event in life’s history falls into the realm of contingency” (1989, p. 290), so that “any replay of the tape would lead evolution down a radically different pathway from the road actually taken” (1989, p. 51). The consequence is that all forms of life might have forever remained simple. But there is a flaw in this reasoning. Contingent historical processes can be governed by laws, as the second law of thermodynamics proves. So the fact (and it presumably is a fact) that life is a contingent historical process is entirely consistent with the arrow hypothesis. The sign of this would be that replaying the tape many times and pooling the results would reveal a robust pattern of increasing maximal complexity. Even though the details from different runs would vary due to different historical accidents, the overall pattern of increasing maximal complexity would still obtain. How could Goald have missed this? Perhaps he did not actually attempt to replay the tape of life but instead was content just to imagine what would happen.

In *Darwin’s Dangerous Idea* (1995), D. C. Dennett is confident that the arrow of complexity hypothesis is true. On analogy with so-called “forced moves” that are obvious to any decent chess player, Dennett argues that complexities such as tools and language offer such obvious adaptive advantage that the stochastic process of natural selection would almost inevitably discover them. The forced move mechanism is a relevant consideration, and under certain conditions it would predict a robust regularity of increasing maximal complexity of life. But it is unclear whether forced evolutionary moves really do vindicate the arrow hypothesis. For one thing, it is unclear whether the relevant forms of complexity really do offer an appropriately large adaptive benefit over the alternatives. After all, simplicity provides its own benefits, and the final weighing of alternatives is highly sensitive to environmental context. How much confidence in the advantages of complexity is really warranted? Furthermore, even if complexity were advantageous, it is unclear whether natural selection really could discover them in a feasible amount of time, or even at all. Even if we can imagine adaptations that we could prove on engineering grounds would be clear adaptive winners, there is no guarantee that

evolution can discover them by simple genetic manipulations of existing forms of life. Not even natural selection can turn every sow's ear into a silk purse. How could Dennett have been so confident in the arrow of complexity? Perhaps like Gould he did not actually attempt to replay the tape but only imagined what would happen.

Other biologists and philosophers have said things that would question or explain the arrow of complexity hypothesis, but in general they are just verbal speculations like those of Gould and Dennett. It is fine to imagine what would happen if the tape of life were replayed, provided one does not place excessive weight on the results. Complex systems are inherently unpredictable (e.g., Wolfram 1994), so imagination is notoriously faulty about their robust properties. Furthermore, the evolving biosphere is about as complex a system as there is. A central lesson of the study of complex systems is that their emergent properties can usually be discovered only by actually observing them.

But how could we actually replay the tape of life? After all, we have evidence about just one evolution of the biosphere (barring extraterrestrial travel and wonderful luck) and we cannot actually replay that tape. Now and then nature provides surrogates for replaying the tape, as when the biosphere recovers after a mass extinction, or when an isolated island is newly created and life first populates it, or when we discover a complex ecosystem that has been largely isolated from the rest of the biosphere. But there are obvious problems with these natural experiments. Too many pre-existing life forms survive the mass extinctions and contaminate the "new" biosphere, isolated islands are repopulated by existing, highly evolved forms of life, and the isolated ecosystems are not isolated enough. The obvious solution is to replay a mathematical model of an evolving system. If the model is simple enough, we can calculate its generic properties *a priori*. If the model is at all complex, we simply resort to the next best thing and simulate it on a computer. Observing the effects of smoothly sweeping through parameter space then reveals the model's generic behavior. This procedure produces hard evidence rather than wishful thinking. "Putting your model where your mouth is" in this way replaces verbal speculation with computer simulations and enforces reference only to explicit and feasible mechanisms (Bedau 1998).

Complexity growth by passive diffusion

Dan McShea's models of random phylogenies (McShea 1994, 1996; see also Gould 1996) contain a population of species, and each species has a number representing the complexity of the organisms in the species. The complexity of a daughter species is like its parental complexity, but mutations change daughter species complexity. Over time, extinction and speciation events change the composition of the species in the population, and the population evolves. A key

feature of these models is that speciation and extinction events are completely random and unbiased with respect to complexity. The species that go extinct or speciate are selected from the population irrespective of their complexity. A species' complexity does not univocally affect its probability of speciating or going extinct. As a result, the model biosphere passively diffuses through the space of possible species. The models ordinarily assume that the earliest species were among the simplest possible and that there is a "left wall" of minimal complexity for any possible organism. Given these constraints (and various assumptions about rates of speciation and extinction), one can generate random phylogenies and thus replay the tape to one's heart's content. Figure 1 shows six examples of replaying such a model.

The first thing to appreciate is that this model seems to vindicate the arrow of complexity hypothesis. That is, a robust regularity of this model is that the complexity of the most complex organisms increases over time. There can be exceptions, but they tend to be transient. The mechanism behind this robust regularity is simply diffusion in the space of recipes. No matter what the complexity distribution of the initial species is, random speciation will robustly generate more and more complex species over time.

McShea himself emphasizes a different point: since a trend of increasing maximal organism complexity could be due simply to passive diffusion, a trend of increasing complexity does not show that complexity provides any adaptive advantage in general. Gould echoes this point in *Full House* (1996), a successor to *Wonderful Life* (1989). My concern here is with a slightly different question: What if anything do passive diffusion models show about the arrow of complexity hypothesis?

I would argue that this kind of model fails to explain the trend of increasing maximal complexity. To appreciate this, notice that a species's "complexity" in this model is entirely nominal. That is, a species does not have any real features that actually have any real degree of complexity. Setting the intended interpretation aside and simply looking inside the model itself, each species has the same complexity—they are each simple data structures with one numerical property. That property can be interpreted as complexity, but it could equally well be interpreted as any other scalar value. For example, the property could be interpreted as the species's degree of intelligence, or as the mean mass of its adult members, or as its entire biomass, or as the mean number of cell types of its members, or as its geographical range, or as the number of limbs of its members. Thus, the model does not really explain differences in complexity. It *presumes* an indefinitely large space of possible species that is structured by a relative complexity relation and that affords unlimited accessible complexity increases. It is true that diffusion will tend to increase maximal complexity given such a space, but this begs the question of whether this space exists, what constructs and structures it, and how the process of evolution enlarges those sub-regions that are accessible to

further evolution. Without answers to those questions, any explanation of how complexity increases is excessively shallow.

Here is another way to see the problem. If these models really explained how complexity increases, then reinterpreting “complexity” as some other property P would produce an equally good explanation of the robust regularity of increasing maximal P in the evolution of life. To take a few extreme examples, we could interpret P as intelligence and have an argument for the inevitable increasing maximal intelligence. Or we could interpret P as moral sensitivity and explain inevitable moral improvement as a robust regularity. Or we could relabel the X-axis as the number of major evolutionary transitions that have occurred and explain the major transitions in evolution as a result of passive diffusion. Obviously, any such explanations would be excessively shallow at best. Even if such arrows existed (and it’s not clear that they do), these “explanations” beg all the interesting questions and make the evolutionary robustness of a trend all depend on the prior presupposition that the right space of accessible evolutionary possibilities exists. When and only when such spaces exist will passive diffusion produce robust increasing trends. So the critical issue is whether and why such spaces exist and can be captured in a model. Passive diffusion models simply assume the affirmative.

Passive diffusion models do make one important point: If there is a predefined genotypic space of organisms of varying “complexity,” then passive diffusion can passively explore all accessible regions of all dimensions of this space. So, the problem of modeling a trend like the arrow of complexity reduces to the problem of creating the appropriate predefined genotype space. That problem remains unsolved.

The evolution of complexity in artificial life models

This problem would be avoided with a richer model containing entities with properties with the right intrinsic meaning. What would such a model look like? If one is interested in the arrow of complexity hypothesis, one might ask for an evolving population of entities with intrinsic properties that locate them in an infinite space of genetic possibilities and which give them some natural intrinsic complexity. One would like the population’s exploration of genetic space to be governed by natural selection operating with an endogenous fitness function that dynamically and unpredictably changed as the entities co-evolve. One would like the organisms’ to co-construct their effective environment, rather than reacting a pre-existing and fixed environment. Since this is a computer model, one could even ask that the system itself be computationally universal, in the sense that any precisely describable process could be modeled in the system, in order to not to limit the system’s evolutionary potential. Many people have thought that the

combination of these features should be sufficient to generate a trend of increasing maximal complexity. The good news is that a number of such models are studied in “soft” artificial life (“soft” because it is based in software, by contrast with “hard” artificial life which concerns robots and other hardware, and “wet” artificial life which occurs in a wet lab; see Bedau 2003a). These are *constructive* models in the sense that they involve creating a system that actually possesses, rather than merely can be interpreted as possessing, the properties of interest. They are also *emergent* models in the sense that, even though we have complete knowledge of the models’ underlying rules and their initial and boundary conditions, it is impossible to predict or derive their global behavior—such as any trends concerning organism complexity—without observing their behavior unfold over time (Bedau 1997).

The most famous artificial life model with all of the above properties is Tierra (Ray 1992). This consists of a population of self-replicating computer programs that compete for space and CPU time in computer memory. They are loosely analogous to an evolving population of RNA strings in a chemostat. The program is typically seeded with an ancestral self-replicating program that is designed by hand. (Tom Ray’s original ancestor was 90 lines of computer code.) Programs are each allocated CPU cycles, and they eventually make copies of themselves in memory. When memory comes close to getting full, older programs are removed from the system to make room for new programs. Random errors (“mutations”) are introduced into programs from time to time, so that the population of program varies. Most mutations break a program so that it is no longer able to make copies of itself, but once in a while a mutation is beneficial and allows a program to reproduce more quickly. This program will tend to produce children faster than others and so will spread through the population. Tierra became famous because a complex ecology of different kinds of programs evolves spontaneously, without having been anticipated or designed in advance by the programmer. This ecology includes parasites that take advantage of the copy instructions in their neighbors, hyper-parasites that evolve a defense against such parasites, mutualists that reproduce only in a spatially-localized colony of similar programs, and cheaters that exploit such colonies.

For present purposes, it is critical to appreciate that Tierran “creatures” (i.e., self-replicating programs) have an intrinsic rather than merely nominal complexity. For example, the 90-instruction ancestral program has an explicit copy loop that the 40-instruction parasites lack. Even without a precise definition of complexity, it is obvious that the ancestor creature is more complex than the parasite creature. It is also critical for present purposes to appreciate that Tierra does not beg but squarely faces the right questions. It has an actual space of evolutionary possibilities—the infinite universe of programs that can be written in the Tierran computer code. The actual population of Tierran programs residing in memory at any given time inhabits a small subset of that space. This space of possibilities has an intrinsic structure with respect to complexity; different programs are intrinsically more or

less complex than others. Moreover, the range of possible mutations determines which nearby regions of genetic space are accessible to others. Evolution changes the accessible subspaces by changing what subspace is currently inhabited. E.g., once the population has discovered a certain kind of adaptation such as unrolling the loop (Bedau and Brown 1999), further variants on that adaptation are accessible. More interestingly, since some kinds of creatures depend on interactions with others, some regions of genetic space cannot be inhabited unless certain other regions are simultaneously inhabited. For example, parasites cannot survive without hosts to exploit, and cheaters cannot survive without mutualists to exploit. Thus, the Tierra model allows us to study when, why, and how evolution creates complex “creatures.”

So, how creative is evolution in Tierra? In particular, does Tierra vindicate the arrow of complexity hypothesis? The short answer is “No.” Despite the complex ecology that spontaneously arises, as far as anyone knows Tierra fairly quickly stops producing anything new. The players in the ecology described above might change, but the population always consists of essentially the same kind of creatures. The production of qualitative novelty peters out below a complexity ceiling, contrary to the arrow of complexity hypothesis. This behavior of Tierra does not disprove the arrow hypothesis, of course. Rather, it suggests that Tierra is an inappropriate model for testing the hypothesis. Tierra fails to produce a trend of increasing complexity even though it encompasses creatures that co-construct their environment and evolve in an infinite genetic space by natural selection with an endogenous fitness function fluctuating via co-evolutionary dynamics, and even though the Tierran language is computationally universal. So, Tierra is a *counterexample* to the sufficiency of those mechanisms for producing an increasing maximal complexity trend. Maynard Smith and Szathmáry (1995) have said that the theory of natural selection does not predict that organisms will get more complex. Study of Tierra shows that the same can be said when we add creatures that construct their environment while evolving in an infinite genetic space according to endogenous fitness criteria that are unpredictably buffeted by the process of co-evolution. This is negative news, but it is constructive progress because it motivates the search for new and better models.

Tierra is not an isolated example. One problem or another besets all other attempts to use contemporary models test the arrow of complexity hypothesis. For example, consider John Holland’s Echo model (Holland 1992, 1995). An Echo world consists of a collection of sites, with agents migrating from site to site. The agents need resources to survive, and their activity is all geared around continually finding more of these necessary resources. Different letters of the alphabet represent different types of resources available in the world. A fixed amount of resources is distributed to each site at each time step. Each Echo agent’s genetic material (its “chromosome”) is composed out of the resources. The chromosome has eleven sub-strings; these constitute an agent’s “tags” and “conditions” together with a

genetic specification of which resources the agent can ingest from the environment. An agent's tags are "external" in the sense that other agents have access to them, while an agent's conditions are "internal" in that they are inaccessible to other agents. The tags and conditions are used to determine the outcome of the three types of interactions that Echo agents can engage in—combat, trade, and mating. Whether two agents interact and what type of interaction they have is determined by comparing the agents' tags and conditions. External tags and internal conditions allow complex non-transitive relationships to exist. Agents gain resources from their environment and lose them through a metabolic tax. They also can gain and lose resources by interacting with other agents.

Holland's Echo model contains a number of plausible ideas, perhaps most centrally the suggestion that evolution fundamentally affects how agents in a population interact, and it was designed explicitly to illustrate the type of complex adaptive system that could exhibit spontaneous complexity growth (Holland 1995). So it would be natural for someone to suspect that Echo would show the growth of complexity. However, as it happens, the version of Echo that is implemented today (Hraber et al. 1997) fails to show any signs of this kind of behavior (Smith and Bedau 2000). It is worth pointing out that this version of Echo is a simplified model that fails to implement some of Holland's original suggestions, so it might turn out that Holland's original hypotheses about complexity growth are still true. However, since the behavior of these complex adaptive systems is emergent, there is no way to tell whether Holland's suggestions are true until someone implements them and then empirically studies the resulting model.

Another model that has produced some impressive behavior is Avida. This model is based on Tierra but includes some interesting modifications, most notably the following three: (i) The space in which creatures interact is two-dimensional, which is more realistic than the one-dimensional spatial structure in Tierra. (ii) Creatures in Avida are typically prevented from reading the code of their neighbors, so the informational parasitism that drove much of Tierra's ecological dynamics is blocked. (iii) Creatures in Avida have the ability to input and output numbers and can garner extra energy from their environment by using these information channels to perform certain logic functions. Replaying the tape with Avida has produced some quite interesting results. Notable among these is an explanation of the evolutionary origin of complex organismal features (Lenski et al. 2003). This study demonstrated in concrete detail how natural selection can create complex features by modifying existing structures and functions.

Although the authors of this study were not intending to address the arrow of complexity, their results might nevertheless seem to illuminate our question. However, this is not the case, for a variety of reasons. For one thing, the complex features (certain logic functions) that evolved in Avida were explicitly anticipated by the model's designer, and the steps toward their creation were explicitly

rewarded. Moreover, the “complex” features are certain logic functions that can be created out of simpler logic functions, and earlier populations were explicitly seeded with those simpler logic functions. So this result does not qualify as the kind of qualitative complexity growth that we are seeking to understand. Finally, as with Tierra, the complexity growth in Avida always reaches a plateau and stops. Thus, these results with Tierra do not capture the kind continual growth in qualitative complexity that we can observe in the biosphere, which is what the arrow of complexity hypothesis seeks to explain.

There is another study using Avida that apparently does directly address the arrow of complexity hypothesis (Adami et al. 2000). The authors investigate whether it is possible to make a case for a trend in the evolution of complexity in biological evolution, and in part because it is tractable, they focus on genomic complexity, which is reflected in the number of essential loci in a genome (i.e., the loci that would be lethal to alter). After replaying the tape numerous times, the authors conclude that genomic complexity is forced to increase, as the genome stores increasing amounts of information about the organism’s environment. Interesting as these results are, though, they do not actually shed light on the arrow of complexity. As with the earlier Avida result, there is no increase in qualitative complexity. The genetic possibilities are predefined and static, so natural selection is simply fixing existing loci. Furthermore, the system is simply adapting to a pre-existing environmental complexity. The environment is fixed throughout the course of these experiments, so they do not explain how evolution creates new environmental complexity. And finally, as before, what complexity growth there is soon stops. So once again, Avida does not capture the kind of continual growth in qualitative complexity that we seek to understand.

There are a number of other interesting developments that bear on our question about the arrow of complexity hypothesis. Some developments are verbal or analytical mathematical theories, and others are constructive models of biological evolution. My remarks about them would be like those above. My general response to the merely verbal proposals is that they will remain inconclusive until they follow the maxim to “put your model where your mouth is.” On the other hand, the mathematical and constructive computer models of evolution that I have seen to date all tend to suffer fates similar to those we encountered with passive diffusion models, Tierra, Echo, or Avida (Bedau et al. 2000). The key point is that none has yet produced a convincing example of the growth of organism complexity.

My diagnosis is that present models do not allow evolution continually to create opportunities for qualitatively new kinds of adaptations. Sterelny and Griffiths (1999), following Maynard Smith and Szathmáry (1995), suggest that the mechanisms of evolution have changed in ways that open up new evolutionary possibilities. What is missing in present models, I think, is the opportunity for the

mechanisms of evolution to change in this way—and especially for it to happen not by some *ad hoc* device but as a result of the process of evolution itself. This is not impossible in principle, but it is missing in practice in today's models. The key question, then, is to determine what mechanisms would suffice to achieve this kind of qualitatively expanding evolutionary process.³ I should emphasize that my diagnosis is just a hunch. People used to have other hunches, e.g., that an infinite genetic space explored by a dynamic endogenous fitness function would achieve significant evolutionary creativity (Packard 1989), but Tierra showed that such hunches were false. My hunch might suffer the same fate in the end.

The bottom line contains both good and bad news. The bad news is that today we have no model that produces a trend of ever increasing maximal organism complexity. No existing model is sufficiently creative to even test the arrow of complexity hypothesis. The good news is that soft artificial life provides a straightforward research program of constructively testing new proposed mechanisms (like the one in my diagnosis above) with new computer models.⁴ I don't want to give the misimpression that this research program is easy. It can be difficult to construct a model that captures one's hunches about the key mechanisms behind evolutionary creativity. A perfect example of this is my own hunch described above, for I have not yet succeeded in capturing that hunch in a model. The proper conclusion to draw from this is that I do not fully understand the mechanism proposed in my own diagnosis above. For if I did, I would immediately make a model in which evolution continually creates opportunities for qualitatively new kinds of adaptations. But I have not yet devised a way to achieve this in a

³ No doubt niche construction (Odling-Smee et al. 2003) will figure significantly among the mechanisms in a model of a qualitatively expanding evolutionary process, but it is important to point out that this mechanism is far from sufficient. For Tierra and most other artificial life models have organisms that can continually reconstruct their own environment as they evolve, and thus can continually modify the very environment to which they are adapting. Yet these models are not qualitatively expanding evolutionary processes.

⁴ Peter Godfrey-Smith once used artificial life to illustrate the idea of a degenerating scientific research program (Godfrey-Smith 2003). Since I am arguing that artificial life models are the best way to explain the arrow of complexity, I should respond to Godfrey-Smith's suggestion. I have two main points. First, in point of fact there has been and still is plenty of exciting progress on many fronts in artificial life (see, e.g., the recent review in Bedau 2003a). At least that is the general consensus of those who are following the latest highlights. Second, those in artificial life specifically trying to explain the evolution of complexity view it as one of a small number of grand challenges, of the sort that could take on the order of 10-25 years or more to solve (Bedau et al. 2000). So lack of quick success is not good evidence that this grand challenge will not be met in the near future; that can be determined only after we have the fruits of another generation of sustained effort.

model. My failure to date illustrates my earlier point about the importance of moving beyond mere words and putting your model where your mouth is. It is only if and when I succeed in putting my model where my own mouth is that we will learn whether my diagnosis above has real merit—or even what exactly my diagnosis really means.

Objections and replies

My thesis about the trend of increasing complexity and the arrow of complexity stands in need of some clarifications and qualifications. A good way to develop them is by replying to objections that might be raised against my foregoing argument.

Self-sealing argument. My argument might seem self-sealing for the following reasons. I argued that Tierra was an inadequate model for our purposes on the grounds that it did not produce a trend of increasing organism complexity. Any model adequate to test the arrow hypothesis must be able to produce this trend. But then won't the model necessarily confirm the arrow hypothesis? If a model fails to confirm the hypothesis, won't it be ruled inadequate like Tierra? If so, then it is impossible to disprove the arrow hypothesis and my argument is self-sealing.

This worry is based on a misunderstanding of replaying the tape of life. The issue is whether the observed increase in complexity is merely an accident, or whether it is a robust feature of systems like the biosphere. Since we cannot actually replay the course of evolution, we seek to do the next best thing and replay a model of reality. When looking for adequate models to test the arrow hypothesis, we must consider three questions. The first is whether the model can produce a trend of increasing maximal organism complexity, the second is whether the model is like reality in the relevant respects, and the third is whether the model produces the trend robustly or merely as an accident. If like Tierra a model cannot exhibit a trend of increasing complexity, then the model should not be used to replay the tape. The biosphere *does* produce the trend, so the model's inability to reproduce that behavior disqualifies it as an adequate model of the biosphere. On the other hand, if the model can produce the trend, then it remains a candidate for replaying the tape.

Our next question is whether the model is like the biosphere in the relevant respects. Since the point of experimenting with models is to discover what would happen if we could replay the course of evolution, the models must capture the key features of reality. For example, they should presumably exhibit the process of natural selection operating on an indefinitely large space of genetic possibilities. Of course, the models will simplify away a great number of properties of the actual biosphere, so they will purposely be unrealistic in those respects. Exactly which

features of reality are crucial to capture in an adequate model is one of the open questions that we hope to answer by our attempts to replay the tape, and the proper answer to this question will be disputed (see below). But there are still plenty of features that will clearly disqualify a model. For example, we should disqualify any model that explicitly gives an *ad hoc* fitness boost to organisms that is proportional to their complexity, for this model would beg the very question under investigation. (Some versions of Avida use a fitness function that explicitly rewards organism complexity, as we noted above.) Similarly, we should presumably dismiss any model that assumes that all complexity increases are simply random mutations on a par with all other random mutations (as simple random diffusion models do).

If a model can produce the trend of increasing complexity and it seems realistic in the relevant respects, then it is a plausible candidate for replaying the tape. The key question now is whether the trend of increasing complexity is a robust feature of the model. Note that this is an open question. There is no guarantee that a model that can produce the trend will do so typically or generically. Those like Gould in *Wonderful Life* who think that the observed trend of increasing complexity is mere contingency should expect the same kind of behavior from the model. It is entirely possible that when the model produces a trend of increasing complexity, this is only an accident. Thus, my argument is not self-sealing. What matters is whether realistic models that can produce a trend of increasing complexity do so typically.

For the sake of simplicity, I am purposely downplaying certain complications that might arise when replaying the tape. Eventually we might have two realistic models that can produce the trend, but the trend might be robust only in one. In this and similar circumstances, we would need to take a closer look at what explains the difference between the models and how the models compare with what we know about reality. We might then choose to frame more detailed hypotheses and test them. All of this could be done in a constructive and empirically grounded manner. But first things first. Before we worry about this hypothetical situation, we must deal our actual situation and find any relevantly realistic model that produces a trend of increasing organism complexity.

Saved by computational universality. A number of artificial life models like Tierra are computationally universal (Maley 1994); that is, they have the power of a universal Turing machine and can be programmed to compute any algorithm or simulate any mechanical process. For any such model, it might be thought to be trivially, even mathematically, true that the model can produce a trend of increasing complexity. For surely there is a mechanical process that produces such a trend; after all, the biosphere is exhibiting the trend and it surely is a mechanical process, at bottom. So any computationally universal system would be capable of producing that trend. And thus the arrow of complexity hypothesis would be trivially true.

There are two reasons we can safely dismiss this worry. The first we canvassed above. Even if a system can produce a trend of increasing complexity, it is still an open question whether that trend is robust. So, far from being trivially true, the arrow of complexity hypothesis might still be false for the model in question. The second reason not to think that computational universality trivializes our question involves distinguishing what is possible in principle from what is possible in practice. The argument from computational universality shows only that producing the trend of increasing complexity is possible in principle. But what is relevant here is a model's ability to produce the trend in practice. The biosphere is actually producing the trend, so obviously it can produce it in practice. Thus, a model that can produce the trend merely in principle but not in practice is unlike the biosphere in the relevant respect and so is unsuitable for replying the tape. And of course, plenty of things that are possible in principle are impossible in practice. So computational universality is a red herring. It is not sufficient for showing that a model will vindicate the arrow of complexity hypothesis, nor even for showing that a model can in practice exhibit a complexity growth trend.

Conway's famous Game of Life (Gardner 1983, Poundstone 1985) nicely illustrates the irrelevance of what is possible merely in principle. This "game" is "played" on a two-dimensional rectangular grid of cells, such as a checker board. Time is discrete. A cell's state at a given time is determined by the states of its eight neighboring cells at the preceding moment, according to the following birth-death rule: A dead cell becomes alive if and only if 3 neighbors were just alive, and a living cell dies if and only if fewer than 2 or more than 3 neighbors were just alive. (Living cells with fewer than two living neighbors die of "loneliness," those with more than three living neighbors die of "overcrowding," and a dead cell becomes alive if it has the three living neighbors needed to "breed" a new living cell.) The Game of Life is a cellular automaton—the sort of system that Wolfram has promoted as a paradigm of complex systems (Wolfram 1984, 2002).

Some Game of Life initial configurations remain unchanging forever, others oscillate indefinitely, still others continue to change and grow indefinitely. One simple and striking example—dubbed the "glider"—is a pattern of five living cells that cycles through four phases, in the process moving one cell diagonally across the Life field every four time steps. Some other notable patterns are "glider guns"—configuration that periodically emit a new glider—and "eaters"—configurations that destroy any gliders that collide with them. Clusters of glider guns and eaters can function in concert just like AND, OR, NOT, and other logic gates, and these gates can be connected into complicated switching circuits. In fact, Conway proved (Berlekamp et al. 1982) that these gates can even be cunningly arranged so that they constitute a universal Turing machine, and hence are able to compute literally every possible algorithm, or, as Poundstone vividly puts it, to "model every precisely definable aspect of the real world" (Poundstone 1985, p. 25). So, since

the trend of increasing complexity is a “precisely definable aspect of the real world,” the Game of Life can model that trend—at least, in principle. However, there is no reason to think that it can model that trend in practice. The problem is that the Game of Life configuration that constitutes a universal Turing machine would be so large, and the algorithms for modeling the trend would be so slow, that the space and time requirements for the computation make the task completely impractical.⁵

Enough space and time. At this juncture a related worry arises: Maybe we have failed to observed Tierra and similar models produce a trend of increasing complexity merely because we have failed to give the models enough memory space or enough CPU time to show what they can do. Maybe we need only give them a larger population or let them evolve longer. When one compares population sizes and numbers of generations, the space and time resources available for existing models seems woefully small compared to the biosphere. A typical artificial life model has a population size between 100 and 10,000 organisms, with the largest population sizes being on the order of 1,000,000. By contrast, the number of organisms in the biosphere is many orders of magnitude larger. Similarly, life has existed on the Earth for billions of years, and over 500 million years have elapsed since the Cambrian explosion. Even with the fastest computers today, artificial life models have not produced anything approaching these numbers of successive generations. Thus one might conclude that Tierra and other models just need comparable numbers of organisms undergoing comparable number of generations in order to exhibit comparable trends in complexity growth.

⁵ A brief document available on the web reports the invention of a (non-universal) Turing machine implemented in the Game of Life (Rendell 2001). The device employs 1714 x 1647 cells (and uses additional cells to represent the tape). It uses 11,040 Game of Life generations to accomplish one Turing machine cycle. I am not aware whether this report has been independently verified. It is reported that the Turing machine had 3 states and 3 symbols, but that the methods could be extended to produce a 16 state and 8 symbol universal Turing machine. Such a universal Turing machine would reportedly be almost half again as slow, taking 15,360 Game of Life generations per Turing machine cycle. Representing the universal Turing machine would take more than an order of magnitude more cells (about 10^7), and storing even a very simple program for the universal Turing machine would add increase the size of the initial configuration by about a couple orders of magnitude (10^9). A program that would implement an evolutionary process exhibiting a trend of increasing organism complexity would be many orders of magnitude larger than this. These numbers indicate the sense in which it would be completely impractical to model the trend of increasing complexity in the Game of Life, even though this is possible in principle.

One way to make this worry vivid is to consider the following simple thought experiment involving complexity growth. Imagine a Petri dish covered with nutrient agar that is dotted with a number of bacterial colonies. You can also imagine continually inducing mutations in the bacteria. Now, ask whether you would expect to observe a trend of increasing qualitative complexity. Of course, you would not see any such trend. The bacteria might evolve a bit because of the mutations, but they would remain bacteria, and indeed essentially the same kind of bacteria. But now notice that this imaginary system is much more complex than any artificial life computer model; the Petri dish contains vastly more organisms, each one of which is itself a self-assembling and self-organizing molecular machine that is vastly more complex than the most complicated cell simulation. So, if the Petri dish system is too simple to exhibit complexity growth, how could we possibly expect it of any existing artificial life model—unless its spatial and temporal resources are vastly increased?

The first point to make in reply to this worry concerns the Petri dish thought experiment. There is an obvious sense in which the Petri dish system would be more complex than any artificial life model, in the sense that the bacterial population is orders of magnitude larger than the model population and number of molecules in each bacterium is order of magnitude larger than the number of bits used to represent each virtual organism. But that might not be the relevant comparison, for those might not be the properties that matter. Once population size is large enough that random genetic drift does not dominate natural selection, it is not evident that population size itself directly affects the existence of trends concerning the evolution of organism complexity. Similarly, it is not at all evident that the number of molecules in an organism reflects the complexity of the kinds of behaviors or interactions which the organism can engage—yet this latter is presumably the kind of property that bears on whether a system will exhibit a trend of increasing organism complexity. Another property that presumably would affect such trends is the complexity of a system's niche structure: how complex each niche is, and how complexly they are all connected. But an artificial life model might well be more complex than the bacterial population in the Petri dish with respect to the properties that matter. It is relatively easy to engineer models in which virtual organisms can engage in a very large number of different kinds of behavior and interactions. Similarly, while the bacteria in the Petri dish would all experience essentially the same kind of environment, it is easy to engineer a model with an environment that has very different kinds of niches connected in a very complex topology. Thus, it is not at all clear that the Petri dish system is more complex than artificial life models in the relevant respects.

The main point to make in response to the conjecture that existing models would behave qualitatively differently if given vastly more space or time is that we simply have no compelling evidence to support this conjecture. What would such evidence look like? One obvious form of evidence would come from experiments

in which spatial and temporal resources are increased incrementally. If such increases often or even sometimes did cause the models to evolve qualitatively more complex organisms, this would certainly count in favor of the conjecture. However, I am unaware of any such experiments having been done thoroughly and carefully. And the anecdotal experience of myself and my colleagues all points in the opposite direction. When we increase the space or time available to Tierra and other models, they continue to exhibit the same kind of behavior and qualitatively more complex organisms never appear. Of course, this does not prove that Tierra or other existing models cannot produce a trend of increasing organism complexity given yet *more* space and time. Today this remains an open empirical issue. It is worth emphasizing that this issue can be addressed constructively and creatively by devising and experimenting with artificial life models.

Hopelessly complicated. One might conclude from my argument not just that Tierra is missing some key mechanism and that we failed to provide it with sufficient time and space, but that the model itself is much too simple. Actual biological organisms are vastly more complicated than any virtual organism in a model, and the web of interactions that connect actual organisms in the biosphere is vastly more complicated than the networks connecting virtual organisms in a model. As biologists learn more and more of these details, we cannot help being struck by how complicated the biosphere is. No doubt many of these complications figure into the process that produces the trend of increasing organism complexity. On the other hand, there is a limit to how complicated any practical computer model can be, because of constraints on the number of lines of code that human can write and understand. One might conclude from this that no practical computer model could ever be complicated enough to capture the features that enable the biosphere to exhibit the complexity trend. Too many details about the biosphere are critical, and any sufficiently realistic model would need to be hopelessly complicated.

If this worry is sound, then we will never devise a model that is adequate to replay the tape of life. If I am right that we must replay the tape in order to evaluate the arrow of complexity hypothesis, then this worry entails that the arrow of complexity hypothesis will remain shrouded in mystery. We will never learn whether the remarkable growth in complexity that we can observe in the biosphere is merely an accident or an inherent property of evolving systems like the biosphere. However, there is no reason yet to draw this bleak conclusion. We have no good reason to think there is no simple feasible mechanism that will produce the complexity trend. Plenty of simple mechanisms have not yet been explored. In addition, plenty of more complicated models could be devised. So it is much too premature to conclude that any adequate model must be hopelessly complicated. It is another open empirical question—and one that artificial life models can again constructively explore—whether practical computer models can capture whatever it is that enables the biosphere to exhibit growing organism complexity.

Try another model. I have been emphasizing Tierra because it embodies six mechanisms—a population of creatures, represented in (i) a computationally universal language, that (ii) co-construct their environment and evolve in (iii) an infinite genetic space by (iv) natural selection with (v) an endogenous fitness function fluctuating via (vi) co-evolutionary dynamics—that many have thought were sufficient to produce a trend of increasing organism complexity. Some people might respond that the scope of my argument is limited to Tierra and that Tierra’s limitations do not show that those mechanisms themselves are incapable of producing a trend of increasing complexity. They might conclude that some other model with the same six mechanisms could well produce the trend, and thus that those mechanisms would be sufficient to explain the trend. If there is a problem, it not with those mechanisms but with Tierra.

This worry fundamentally challenges my argument, and I think it is fundamentally misguided. To see why, we need to be careful about what is at issue. First, I would agree that another model with different mechanisms might well produce a trend of increasing organism complexity, but, of course, that is not the issue. Second, I would also agree that another model with those same six mechanisms *and with some other important difference from Tierra* might well be able to exhibit the trend. But that is also not the issue, for it would not show that those six mechanisms by themselves were sufficient to produce the trend. Rather, it would show that those mechanisms *plus something else* are sufficient for the trend. Since Tierra embodies those mechanisms and yet fails to produce the trend, Tierra shows that those mechanisms together are insufficient to explain the trend. This conclusion is not limited merely to Tierra; it applies to any model that embodies those same mechanisms. I presume that we will eventually devise a model that exhibits a complexity growth trend, and I suspect that it will embody many, perhaps all, of the six mechanisms from Tierra. But it must significantly differ from Tierra; it must contain some further key mechanism.

Some might think that it is trivial to identify the key mechanism missing from Tierra. To them I repeat my earlier injunction: “Put your model where your mouth is.” One of the important morals to draw from current efforts to resolve the arrow of complexity hypothesis is that talk is cheap and inconclusive. As we have seen, it is all too easy to spin a superficially plausible tale about why some mechanism “surely must” yield a trend of increasing complexity. This is exactly what people in artificial life originally thought about the six mechanisms in Tierra—and that is why they created models like Tierra that embodied those mechanisms. But remember that the global behavior of artificial life models is emergent; there is no way to tell how they will behave short of watching them unfold in time. This emergent behavior is the reason why verbal speculation about the arrow of complexity is so inconclusive, and why the acid test of any hypothesis about the arrow is to

implement a model that embodies the hypothesis and study the model's behavior empirically.

The arrow is false. I have been arguing that the growth of complexity in the biosphere has not yet been adequately explained, and that the arrow of complexity encapsulates the central hypothesis to be resolved. Some might be skeptical about this project, because they believe that the arrow of complexity hypothesis is simply false. They are convinced that the biosphere is not exhibiting any robust trend that involves the generic production of increasing complexity. Rather, this is simply an accident. Thus, there is nothing interesting to explain and my project is misguided.

There are two things wrong with this worry. First, I have not been arguing that the arrow of complexity hypothesis is true, but that it is an important open question. For all we know, it might be true and it might be false. At this stage, claims that that the arrow is false are dogmatic; they are just expressions of faith. (The same holds for claims that the arrow is true.) Second, even if the arrow of complexity hypothesis does turn out to be false, there is still something important left to explain: the observed trend in increasing organism complexity. Even if that trend is an accident, it still has an explanation. So, among other things, we should be able to explain how it is possible for that kind of accident to happen.

The fact of the matter is that today we have no compelling explanation for how a trend of increasing organism complexity could happen accidentally. It is not that such an explanation is impossible, but that we don't have one in hand yet. What would constitute such an explanation? The most convincing demonstration would be a model that was arguably like the biosphere in the relevant respects and that could exhibit the trend. If we want to explain how the trend could be an accident, then we should show that the trend is not a generic or typical feature of the model. Rather, the trend should be an accidental feature of the model, but still observable. In this respect, it will differ from passive diffusion models, for they make the growth of complexity a typical trait.

Whether you want to argue for or against the arrow of complexity hypothesis, the critical first step is to produce some appropriate model that actually produces the trend. This will set in motion further steps such as investigating whether the model is like the biosphere in the relevant respects, and assuring ourselves that the model produces real and not merely nominal complexity and does not beg any questions of how the space of genetic possibilities is structured. Without this kind of constructive demonstration, professions about why the arrow of complexity "must" be true or false simply ring hollow.

Odious allies. There is one more objection that I should briefly discuss. I have been arguing that nobody yet knows how the process of evolution could have created a biosphere with extremely complicated organisms from a biosphere containing only

simple organisms. This conclusion should be music to the ears of creationist and intelligent design critics of Darwinism who argue that the process of evolution cannot explain how the biosphere contains complex organisms (see, e.g., Wells 2006, and for a more balanced treatment, Pennock 2001). I have not argued that evolution *cannot* explain complexity, but that it has not yet explained complexity, and the fact that we do not yet know how evolution can create complex organisms of course does not imply that evolution cannot create complex organisms. Further, I have been arguing that the appropriate artificial life models will eventually provide a sound naturalistic explanation for the evolution of complex organisms. But dogged defenders of Darwinism will worry that my argument will give aid and comfort to religious critics of Darwinism, that they will cherry pick what is useful for their purposes and twist it into support for unscientific and anti-naturalistic alternatives to Darwinism.

There are two points to make about this worry. For one thing, it is a bad idea to pretend that we can explain the evolution of complexity before we actually can, for this just supports the claims of Darwinism's critics that the evidence for Darwinism has been "exaggerated," "distorted," or "faked," as claimed on the back cover of Wells (2006). But the main point to make is that it is intellectually dishonest to deny that the evolution of complexity is an open question today. That alone should settle the issue. None of us can prevent the unscrupulous from abusing or misusing what we say, but it would be sad indeed if this fear were to muzzle frank and open investigation of the limits of current scientific explanations. That would simply retard progress towards the more adequate explanations of the evolution of complexity.

Conclusions

A fundamental challenge in the life sciences today is to settle whether the arrow of complexity hypothesis is true and explain the observed trend of increasing maximal organism complexity. Progress on this challenge is slowed in part by the lack of recognition that this issue is still open; many people do not realize that the answer to this important question is still a mystery. Some people think that natural selection given an infinite space of genetic possibilities will inevitably produce more and more complex adaptations. But soft artificial life models like Tierra show conclusively that those mechanisms are in general insufficient to produce a trend of increasing complexity. The proof is simple: The models embody those mechanisms but they don't exhibit the requisite behavior. Mechanisms like natural selection in an infinite space of genetic possibilities might still be necessary for explaining the trend, but Tierra shows that they are not sufficient.

The conclusion to draw from this is that biology needs new concepts, theories, and models if it is to resolve the arrow of complexity hypothesis. The status of the arrow

of complexity remains an open question today. My own hunch is that we are missing some key insight, some important new concept or process or mechanism that will resolve why complexity grows over the course of evolution. Once we discover this new idea, it might seem obvious and almost trivially true, like the idea of natural selection seems to us in hindsight. But also like natural selection, before we discover the new idea, our explanations of the evolution of complexity fail to convince, and we feel like we are groping in the dark.

Fortunately, biology does have a new powerful tool for exploring answers to this question: the constructive models of soft artificial life. These models are essential for making progress on deep questions about complex adaptive systems, such as the arrow of complexity. Because these systems are complex, their global behavior is typically emergent. Thus, the only way to be sure how such a system will behave is to build an emergent model of the system and observe its generic behavior. The process of using constructive models to investigate issues like the arrow of complexity is not fool-proof, however. As we have seen, some models beg the interesting questions, and others fail to produce the appropriate behavior. So, proper use of constructive models requires care and experience. Fortunately, in the right hands, they provide a public, repeatable, and empirically grounded method for making incremental progress on deep questions about biological phenomena. Appreciating how constructive models can resolve the arrow of complexity can start to suggest how similar models can play a central role in answering other fundamental biological challenges in the next decade and beyond.

References

- Adami, C., C. Ofria, T. C. Collier. 2000. Evolution of biological complexity. *Proceedings of the National Academy of Science U.S.A.* 97: 4463-4468.
- Allen, C., M. Bekoff, G. Lauder, eds. 1997. *Nature's purposes: analyses of function and design in biology*. Cambridge: MIT Press.
- Bedau, M. A. 1991. Can biological teleology be naturalized? *The Journal of Philosophy* 88: 647-655.
- Bedau, M. A. 1992. Where's the good in teleology? *Philosophy and Phenomenological Research* 52: 781-806. Reprinted in Colin Allen, Marc Bekoff, George Lauder, eds. 1997. *Nature's purposes: analyses of function and design in biology*, pp. 261-291. Cambridge: MIT Press.

Bedau, M. A. 1997. Weak emergence. In James Tomberlin, ed., *Philosophical perspectives: mind, causation, and world*, vol. 11, pp. 375-399. Oxford: Blackwell Publishers.

Bedau, M. A. 1998. Philosophical content and method of artificial life. In T. W. Bynam and J. H. Moor, eds., *The digital phoenix: how computers are changing philosophy*, pp. 135-152. Oxford: Basil Blackwell.

Bedau, M. A. 2003a. Artificial life: organization, adaptation, and complexity from the bottom up. *Trends in Cognitive Science* 7: 505-512.

Bedau, M. A. 2003b. Downward causation and autonomy in weak emergence. *Principia* 6: 5-50. Reprinted in M. A. Bedau and P. Humphreys, eds. 2007. *Emergence: contemporary readings in philosophy and science*. Cambridge: MIT Press.

Bedau, M. A., and C. T. Brown. 1999. Visualizing evolutionary activity of genotypes, *Artificial Life* 5: 17-35.

Bedau, M. A., J. S. McCaskill, N. H. Packard, S. Rasmussen, C. Adami, D. G. Green, T. Ikegami, K. Kaneko, T. S. Ray. 2000. Open problems in artificial life. *Artificial Life* 6: 363-376.

Berlekamp, E. R., J. H. Conway, and R. K. Guy. 1982. *Winning ways for your mathematical plays*. Vol. 2. New York: Academic Press.

Boogerd, F. C., F. J. Bruggeman, R. C. Richardson, A. Stephan, and , H. V. Westerhoff. (2005). Emergence and its place in nature: a case study of biochemical networks. *Synthese* 145: 131-164.

Dennett, D. C. 1995. *Darwin's dangerous idea: evolution and the meanings of life*. Simon and Schuster: New York.

Edmonds B. 1996. What is complexity? In F. Heylighen, J. Bollen, and A. Riegler, eds., *The evolution of complexity*, pp. 1-16. Dordrecht: Kluwer.

Gardner, M. 1983. *Wheels, life, and other mathematical amusements*. New York: Freeman.

Farmer, D., and A. Belin. 1992. Artificial life: the coming evolution. In Langton, C., Taylor, C., Farmer, J. D., and Rasmussen, S., (eds.), *Artificial life II*, pp. 815-840. Redwood City, CA: Addison Wesley.

- Godfrey-Smith, P. 2003. *Theory and reality: an introduction to the philosophy of science*. Chicago: University of Chicago.
- Gould, S. J. 1989. *Wonderful life: the Burgess shale and the nature of history*. New York: Norton.
- Gould, S. J. 1996. *Full house: the spread of excellence from Plato to Darwin*. New York: Random House.
- Holland, J. H. 1992. *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. Second Edition. Cambridge: MIT Press.
- Holland, J. H. 1995. *Hidden order: how adaptation builds complexity*. Reading, MA: Helix Books.
- Hraber, P. T., T. Jones, and S. Forrest. 1997. The ecology of Echo. *Artificial Life 3*: 165-190.
- Kauffman, S. 2000. *Investigations*. New York: Oxford University Press.
- Lenski, R. E., C. Offria, R. T. Pennock, and C. Adami. 2003. The evolutionary origin of complex features. *Nature 423*: 139-144.
- Lloyd, S. 2001. Measures of complexity: a nonexhaustive list. *IEEE Control Systems Magazine 21*: 7-8.
- Maley, C. C. 1994. The computational completeness of Ray's Tierran assembly language. In C. G. Langton, ed., *Artificial Life III*, pp. 152-159. Reading, MA: Addison-Wesley.
- Maynard Smith, J., and E. Szathmáry. 1995. *The major transitions in evolution*. New York: W. H. Freeman.
- McShea, D. 1991. Complexity and evolution: what everybody knows. *Biology and Philosophy 6*: 303-324.
- McShea, D. 1994. Mechanisms of large-scale evolutionary trends. *Evolution 48*: 1747-1763.
- McShea, D. 1996. Metazoan complexity and evolution: is there a trend? *Evolution 50*: 477-492.

Odling-Smee, F. J., K. N. Laland, and M. W. Feldman. 2003. *Niche construction: the neglected process in evolution*. Princeton: Princeton University Press.

Packard, N. 1989. Intrinsic adaptation in a simple model for evolution. In Langton, C. (ed.), *Artificial life*, pp. 141-155. Redwood City, CA: Addison-Wesley.

Pennock, R. 2001. *Intelligent design creationism and its critics: philosophical, theological, and scientific perspectives*. Cambridge: MIT Press.

Poundstone, W. 1985. *The recursive universe*. Chicago: Contemporary Books.

Ray, T. S. 1992. An approach to the synthesis of life. In C. Langton, C. Taylor, D. Farmer, and S. Rasmussen, eds., *Artificial life II*, pp. 371-408. Redwood City: Addison-Wesley.

Rendell, P. 2001. A Turing machine in Conway's Game of Life. Available on the web at www.cs.ualberta.ca/~bulitko/F02/papers/tm_words.pdf. Accessed November 2006.

Skusa, A., and M. A. Bedau. 2002. Towards a comparison of evolutionary creativity in biological and cultural evolution. In R. Standish, M. A. Bedau, and H. A. Abbass, eds., *Artificial life VIII*, pp. 233-242. Cambridge: MIT Press.

Smith, R. H., and M. A. Bedau. 2000. Is ECHO a complex adaptive system? *Evolutionary Computation* 8: 419-442.

Sterelny, K., and P. E. Griffiths. 1999. *Sex and death: an introduction to the philosophy of biology*. Chicago: University of Chicago Press.

Wells, J. 2006. *The politically incorrect guide to Darwinism and intelligent design*. Washington D.C.: Regnery Publishing.

Wimsatt, W. C. 1997. Aggregativity: Reductive Heuristics for Finding Emergence. *Philosophy of Science* 64: Suppl. 2: S372-S384.

Wolfram, S. 1994. *Cellular automata and complexity*. Reading, Mass.: Addison-Wesley.

Woldfram, S. 2002. *A new kind of science*. Champaign, IL: Wolfram Media.

Figure 1

Caption. Six examples of the phylogenies produced by a random diffusion model of the type discussed by McShea (1994, 1996) and Gould (1996). Time increases from the bottom to the top, and complexity of the organisms in a given species increases to the right. There is a left wall of complexity (complexity is constrained to be positive). Note that in general the complexity of the most complex forms of life inevitably increases over time, due to passive diffusion. That is, if you wait long enough, you will see new increase in maximal organism complexity, from passive diffusion.

