

Evolution Unbound: Releasing the Arrow of Complexity

BAYESIAN INTELLIGENCE TECHNICAL REPORT 2010/2

Kevin B. Korb* and Alan Dorin†
Clayton School of Information Technology
Monash University
Clayton, Victoria 3800
Australia

[The] conclusion [that evolution is a fact], even if well founded, would be unsatisfactory, until it could be shown how the innumerable species inhabiting this world have been modified, so as to acquire that perfection of structure and coadaptation which most justly excites our admiration.

Charles Darwin

Abstract

The common opinion has been that evolution results in the continuing development of more complex forms of life, generally understood as more complex organisms. The arguments supporting that opinion have recently come under scrutiny and been found wanting. Nevertheless, the appearance of increasing complexity remains. So, is there some sense in which evolution does grow complexity? Artificial life simulations have consistently failed to reproduce even the appearance of increasing complexity, which poses a challenge. Simulations, as much as scientific theories, are obligated at least to save the appearances! We suggest a relation between these two problems, understanding biological complexity growth and the failure to model even its appearances. We present a different understanding of that complexity which evolution grows, one that genuinely runs counter to entropy and has thus far eluded proper analysis in information-theoretic terms. This complexity is reflected best in the increase in niches within the biosystem as a whole. Past and current artificial life simulations lack the resources with which to grow niches and so to reproduce evolution's complexity. We propose a more suitable simulation design integrating environments and organisms, allowing old niches to change and new ones to emerge.

Keywords: Biological complexity; diffusion; niche construction; artificial life; open-ended evolution.

*Corresponding author. Email: kbkorb@gmail.com; Phone: +61(03) 9905-5198; Fax: +61(03) 9905-5146

†Email: alan.dorin@infotech.monash.edu.au

1 Evolution's Arrow

A common opinion about evolution has been that it swims against the tide of entropy and in particular that evolution over time constructs more and more complex organisms. Take Jacob Bronowski, for example:

In a history of 3000 million years, evolution has not run backward... Why is this? Why does evolution not run at random hither and thither in time? What is the screw that moves it forward, or at least, what is the ratchet that keeps it from slipping back? Is it possible to have such a mechanism which is not planned? What is the relation that ties evolution to the arrow of time, and makes it a barbed arrow? (Bronowski, 1970)

At its most extreme, this view has evolution progressing from bacteria to invertebrates and thence to vertebrates and mammals and, finally, to the pinnacle of life forms, us.¹ Such a view of Progress, however, ignores some quite basic features of evolution. For example, that the bacteria being “progressed from” still exist today and, indeed, have exactly as long an evolutionary history as we do, since we all have common ancestry. So, progress can hardly be characterized by endurance. Instead, progress has been recast as complexity, and complexity itself has been cast in terms favorable to ourselves; for example, as owning complex neural organizations — an account that fails to address the vast majority of earth's life. An infatuation with ourselves is also behind the paradoxicality of the “C-value paradox”, that our chromosomes appear no more complex than those of other mammals and less complex than those of some plants. The two long-standing antagonists Dawkins and Gould have together, and quite rightly, castigated this view as human chauvinism in an exchange in *Evolution* (Dawkins, 1997; Gould, 1997).

Gould preferred to see in every attempt to characterize complexity and attribute its increase to evolutionary processes this hidden agenda of congratulating ourselves on our own unique wonderfulness. Dawkins, on the other hand, considers the evolutionary increase in complexity to be not just compatible with evolution, but intrinsic to it. Evolution climbs “Mount Improbable” (Dawkins, 1996). Dawkins' line of defence for ongoing complexity increase is to suggest that, whereas adaptive processes responding to the abiotic environment may just track meandering changes in the climate, coevolutionary processes acting between species work to develop coadaptations in trajectories that can be regarded as progressive in an engineering sense. Arms races lead to better weaponry and better defences, including better speed, flight and vision, for example. These coadaptive races are periodically interrupted, to be sure:

Each of these ... arms races powered sequences of evolution which were progressive in my sense. But there was no global progress over the hundreds of millions of years, only a sawtooth succession of small progresses terminated by extinctions. Nonetheless, the ramp phase of each sawtooth was properly and significantly progressive. (Dawkins, 1997, p. 1019)

Here we shall take up this debate. We prefer to leave aside any questions about progress and consider only complexity and its possible increase. In particular, we shall consider Bedau's (2006) *Arrow of Complexity Hypothesis (AOC)*, that evolutionary systems show a systematic tendency to increase the complexity of organisms over time. Its truth for earthly biology *seems* clear:

¹For a skeptical review of this consensus opinion, identifying culprits, see McShea (1991).

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. (Bedau, 2006)

Whereas this is strongly suggested by the fossil record of repeated periods of strong diversification and adaptive radiation following mass extinction events, there is not even the *appearance* of truth for AOC applied to artificial life (A-Life): thus far, no A-Life system has demonstrated anything like the creativity of biological evolution. Yet, the search for open-ended complexity in A-Life has a history reaching back to its very beginnings with John von Neumann (McMullin, 2002; Ruiz-Mirazo et al., 2008), continuing notably with Ray (1990), and remaining an important issue now. This history of continued effort, and failure, to find it has led to its discovery becoming a key challenge for the A-Life community: to build a computer simulation which does demonstrate open-ended evolution (Bedau et al., 2000). This is, of course, a version of Darwin’s challenge above: we must find the How of evolution, beyond knowing that it occurs. It is an especially acute version of that challenge. It is all very well to make up a theory about the How of evolution, but in computer simulation we expect our simulations to *produce the goods*. We consider this to be a hard problem. We will not solve it here, but hope to indicate a path towards a solution.

We shall first examine Gould’s main argument against complexity increase in evolution and find it lacking — that while it succeeds tactically, it fails strategically. We shall briefly consider the empirical evidence in favor of complexity increase and find it at least encouraging for that hypothesis. But, taking Darwin’s admonition above seriously, to find the How of evolution and not just the That of it, we shall go on to consider *how* complexity increase might be achieved systematically by evolution. We present ideas for a new and plausible design for evolutionary A-Life that incorporates ecosystem engineering and niche construction as essential ingredients, allowing for adaptive interactions not just between life’s species, but also between them and their abiotic environments, and so driving adaptive complexity increases well beyond the interspecific coadaptations that Dawkins refers to.

2 Complexity through Passive Diffusion

A simple response to the challenge to account for open-ended evolution can be fabricated using Daniel McShea’s (1994) proposed explanation for natural AOC (adopted by Gould, 1996): that any ecosystem under random mutation will show “diffusion” effects, such that given a non-absorbing barrier on the one side (there is no such thing as negative complexity), species over time must occupy more and more complex parts of the natural design space, if only because of random mutations (see Figure 1).²

Gould makes two minor errors in this formulation. First, it is an application of random walk theory (RWT) to a one-dimensional state space, the dimension being the complexity of the system in a genetic design space. The steps in the random walk are mutations (or other genetic changes) occurring between generations. It does not consider the positions of particles in motion, so it is not literally about diffusion. Nevertheless, in deference to common practice we shall refer to this as diffusion, while drawing upon stochastic process theory to explain it. The second minor error is that that theory delivers a limit probability of one of reaching any given finite distance to the right of the origin of the complexity dimension whether or not there is a

²This explanation builds upon Stanley (1973), who applied the concept of “passive diffusion” to Cope’s rule of size increase.

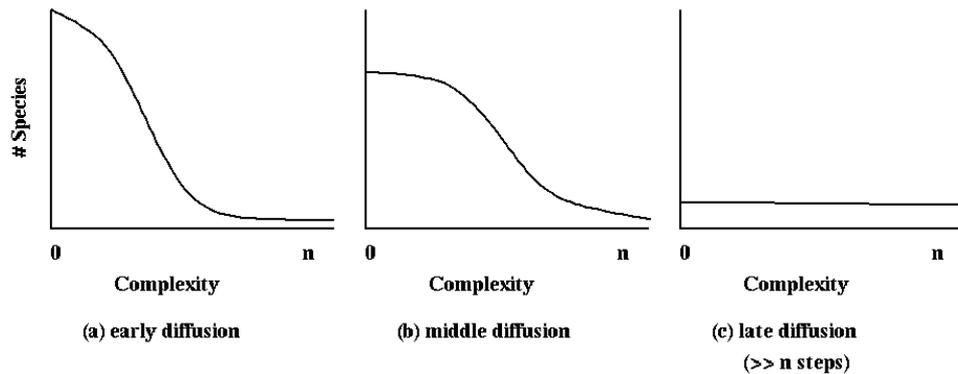


Figure 1: Complexity through diffusion. The unit of the horizontal axis represents complexity changes possible in one generation.

non-absorbing barrier, so long as there is no bias against steps which increase complexity. It is obvious that there is no sense in contemplating negative complexity, so there *is* such a barrier. The barrier is irrelevant, however, to the relevant theorem from stochastic process theory and so to the diffusion explanation (e.g., Ross, 1996, pp. 171-2).³

But the basic idea behind the diffusion argument is right, that random walk theory is already sufficient to explain the evolution of arbitrarily complex species, given no bias (passive diffusion).

Mark Bedau's view is apparently quite different: he claims that the diffusion argument is wrong-headed. Consider Daniel McShea's diffusion simulation of the evolution of species complexity. McShea represents complexity by a number, which is inherited and mutated at random, with equal probability of increasing and decreasing in his passive diffusion experiments, precisely mimicking the requirements of the RWT theorem. In this case, unsurprisingly, McShea's simulation does not contradict the theorem, with species complexity randomly increasing and decreasing in a diffusion across the spectrum of possibilities.

Against the idea that this shows that the evolution of complexity could be explained by passive diffusion, Bedau argues that the concept of complexity here is entirely nominal, or epiphenomal: the complexity number is not *doing* anything in the simulation. And, because of that,

[the] 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' degree of intelligence, or as the mean mass of its adult members, or as its entire biomass, . . . (Bedau, 2006, p. 6)

Or, as has been suggested elsewhere, as its moral virtue or even as *simplicity*! But these last two (unquoted) examples help to reveal an error in the argument. Generally, a good way to show that an argument is invalid is to find an application of its form where the premises are manifestly true and the conclusion manifestly false. But for this to work here there are some clear preconditions: there has to be a clearly definable scalar dimension to substitute for the complexity dimension; that dimension has to be unbounded to the right; there have to be clearly

³Gould (1996) quite explicitly makes this mistake, claiming that the existence of a non-absorbing barrier of minimal complexity is *necessary* for the diffusion argument. This is hardly the only statistical error Gould makes. He considers, for example, the continued presence of a mode of species complexity near the left barrier to be evidence in favor of passive diffusion. In fact, however, such an enduring mode is clear evidence *against* passive diffusion being the sole factor for complexity change, since passive diffusion acting alone results in an approximately uniform distribution over time, as McShea (1994) notes and Figure 1(c) illustrates.

definable mutation operators which can randomly alter species, moving them up or down the dimension without bias. The striking feature of simplicity is that it has a finite upper bound and no finite lower bound, exactly the opposite of complexity, of course. In other words, it is incoherent to try to substitute simplicity for complexity in the argument: given a discrete genetics, it is impossible to forever increase simplicity without striking that non-absorbing barrier of maximal simplicity. As for moral virtue, no one is in a position to clearly define moral superiority, let alone operators on representations of moral virtue which will produce more and less virtuous beings. The absurdity of these examples indeed generates strong intuitions that something has gone wrong, but also reveals that what has gone wrong is a misapplication of the diffusion argument. But the misapplication is already embedded in the earlier examples. In the cases of mean and total mass no sensible simulation, and no physical system, could possibly demonstrate an unbiased random walk. Mass costs energy to build and maintain; there's nothing neutral about it. And intelligence is notoriously difficult to define and measure, so this application of diffusion is thus far strictly imaginary.

So, in rebuttal to Bedau's argument, we claim that, yes, McShea's number can be interpreted as complexity, but it fails to follow that it could be interpreted as simplicity, intelligence, beauty or just any scalar value you care to dream up. There are real constraints on how we can interpret McShea's simulation. In fact, it may well be the case that we *cannot* interpret McShea's number as complexity. In order to do so properly we must first find a clear definition of species complexity, and we must also have a plausible story that increasing complexity of whatever kind we've just defined doesn't impose a fitness cost on the species (negative bias). While we shall offer some ideas about how to measure complexity below, we will not claim to have fulfilled the exact requirements of the RWT theorem, so we will not be explaining away AOC via diffusion. What we do instead is use the diffusion argument to reconsider our target hypothesis.

3 Weak and Strong Arrows of Complexity

We think that what Bedau finds most repugnant in the diffusion argument is that it trivializes the evolution of complexity; in McShea's and Gould's hands (Gould, 1989) it is aimed at undermining the consensus view that evolution runs opposite to entropy. Bedau's underlying intuition, which we share, is that there must be more to the complexity growth of evolution than something as simple as passive diffusion. Any kind of complexity increase that *might* be explained by passive diffusion *alone* is a kind of complexity increase that we are not interested in.⁴ Passive diffusion can be thought of as providing a kind of "neutral shadow" for complexity growth — that complexity that peeks out from the shadows is what is of most interest.

We propose differentiating between two versions of AOC, one which *might* be explained by passive diffusion and another which certainly *cannot* be explained by passive diffusion.

Hypothesis 1 (The Weak Arrow of Complexity) *Evolutionary systems have a robust tendency to produce ever more complex organisms (species). (Bedau, 2006)*

Bedau glosses this by observing that robustness means that the growth in complexity is not merely apparent, but features regularly (robustly, if not universally) when we "replay the tape of

⁴When we say that the diffusion argument trivializes complexity growth, we are *not* saying that diffusion is trivial or unimportant — it is, in fact, essential (cf. Lynch, 2007). Diffusion in evolution is based on mutation, chromosomal cross-over, etc., and is the very stuff of that variation upon which selection acts. So, without diffusion evolution simply stops. What is trivial is its use as a full and complete explanation for complexity growth: there is more to complexity than meets the diffuse eye.

life”, i.e., rerun evolution’s history, as Gould pretended to do (with his imagination; Gould, 1989) and as A-Life researchers in fact do. It is perhaps worth pointing out that this is a thesis about *evolutionary* systems and not non-evolving systems. That is, evolutionary systems are not being supposed to have magical powers, like the ability to overcome the negative impacts of a nearby supernova. The AOC tendency can only be exhibited while the general environment is conducive to the reproduction of the evolved biota. In particular, repeated mass extinctions on earth are no evidence against AOC; in fact, they provide some of the best evidence available that AOC is true (see §5).

Also implicit in the AOC is the idea that the growth in complexity should be the consequence of *how evolution works*; it is *intrinsic* to evolution and not generated by the circumstances in which it operates. Probably, “robustness” is meant to capture this point; however, robust regularities might in fact be due to circumstances within which evolution operates but outside of our control, so that when we “replay the tape of life” the regularities recur, but not due to evolution itself. It follows that if, for example, complexity is constrained by an ecosystem’s energy input, then that is an extrinsic factor that in no way undermines the AOC, so long as evolution indeed reaches that complexity constraint with regularity. Energy and water supply clearly act as extrinsic constraints more local than mass extinction events.

The above remarks apply to both versions of AOC, i.e., including this strong form:

Hypothesis 2 (The Strong Arrow of Complexity) *Evolutionary systems have a robust tendency to become ever more complex.*

Strong AOC is not a minor variant of Weak AOC: they are altogether different beasts.⁵

We focus on Strong AOC for a number of reasons:

1. Weak AOC is subject to a trivializing realization and explanation, in terms of passive diffusion. If the challenge to produce open-ended evolution is a serious one, then it cannot be framed in terms of the Weak AOC, since the challenge of producing a model of passive diffusion is itself not serious.⁶ (See §2.)
2. Strong AOC puts the attention on biosystem complexity and so on such processes as niche construction, ecosystem engineering, and biogeochemical pathways through the ecosystem. These are issues which the A-Life community has largely ignored, and they are potential sources of the creative evolution that has been missing. (See §6.)
3. Strong AOC has an interesting potential relation with entropy, whereas Weak AOC does not. In particular, it is impossible for passive diffusion to run contrary to entropy, whereas whatever might explain Strong AOC must run against entropy. (See §4.)

⁵This point seems not to be widely appreciated. Many discussions segue between talk of organismic complexity, relevant to Weak AOC, into talk of biosystem complexity, relevant to Strong AOC, and back again, without appearing to notice that any transition has occurred. Bedau (2006) simply dismisses Strong AOC as “closely related” to Weak AOC and focuses then on the latter. But our subsequent discussion will make clear that this is incorrect. (For recent treatment of the Weak AOC, see the special issue of *Artificial Life* introduced by Gershenson and Lenaerts, 2008.)

⁶No one has previously produced such a model (other than that of McShea (1994), which suffers from its genome being epiphenomenal). Most A-Life simulations introduce implicit biases against complexity and so cannot be models of passive diffusion. But another impediment is just that no researcher is likely to confuse passive diffusion models with a serious attack on the open-ended evolution challenge in the first place. In any case, for this research we have produced such a model, with a genome that genuinely grows in complexity over time and with that complex genome playing a causal role in the inheritance of complexity over the generations. In consequence, the simulation runs progressively more slowly across those generations. Unsurprisingly, as a proper implementation of an unbiased random walk, our model’s results are qualitatively the same as McShea’s.

3.1 Complexity and Open-ended Evolution

Another contrast we need to make is between the two Arrows of Complexity and open-ended evolution. In a series of papers Mark Bedau and his coauthors introduced a set of statistical measures of the evolutionary activity of systems and used them to rank systems as evolutionarily inactive, bounded and unbounded (Bedau and Packard, 1991; Bedau et al., 1998; Bullock and Bedau, 2006). Unbounded systems, for which there are three criteria of differing strengths offered,⁷ are said to display “open-ended evolution”. Bedau et al. (1998) examined A-Life systems as well as the fossil record for families during the Phanerozoic and concluded that only the latter — i.e., actual evolutionary history — displayed unbounded evolution according to the criteria they developed. This then led to the challenge to the A-Life community to reproduce open-ended evolution (Bedau et al., 2000). Various further investigations and reviews have confirmed this opinion (e.g., Stout and Spector, 2005). Channon (2006), however, in testing his A-Life system Geb has found that it meets the weakest criterion offered for unbounded evolution, that is, it continues to generate new “components” (viable positions in the available genotypic or phenotypic design space) that are useful (adaptive), even while showing bounded biodiversity. What this result establishes, however, is that generating open-ended evolution, at least on this weakest criterion, is not in itself a very impressive achievement. Any system which has a sufficiently large design space and sufficiently benign conditions for evolution can be open-ended in this sense simply by following adaptive gradients in the design space indefinitely. New components will continue to arise, as incremental improvements in adaptive value are always available; and this can be true even without any major evolutionary transitions or the evolution of more complex biological functions beyond those specified in the initial conditions of the simulation. Indeed, the last point is a major defect in Bedau’s evolutionary activity statistics, at least for their employment in assessing the creativity of the evolutionary dynamics: there is no reference whatsoever to the complexity of either the organisms or the ecosystem being evolved. This was inevitable given that Bedau did not attempt to provide a definition of biological complexity, even in his most recent work, which eschews talk of open-ended evolution in favor of talk of the evolution of complexity. Consequently, open-ended evolution under the weakest criterion is more easily realized even than Weak AOC.

So, in sequel we will not refer to Bedau’s classes of unbounded evolution and instead use “unbounded”, “creative” or “open-ended” evolution simply in the sense of satisfying the Strong AOC.

4 Entropy

We now examine the relation between the arrows of complexity and entropy. Information-theoretic entropy is the measure of the expected amount of information required to identify the state (s_i) of a system (S):

Definition 1 (Entropy)

$$H(S) = - \sum_i p(S = s_i) \log p(S = s_i)$$

The second law of thermodynamics states that closed systems increase their physical entropy. At maximum physical entropy we also find maximal information-theoretic entropy. In other words, when a system has minimal macrostructure or high-level organization it requires a

⁷Namely, those with unbounded diversity, or unbounded cumulative adaptations, or both Channon (2006).

maximal amount of information to describe its microstructure, e.g., the location of its particles. Evolutionary systems are physical systems, and they obey the laws of thermodynamics. So, if we are to sustain the intuition that we started with, that evolutionary systems manage to fight entropy, we shall have to reconcile intuition with thermodynamics.

As a clarification, we note that increasing complexity and minimizing entropy (the most extreme form of “fighting” entropy!) are not the very same thing. Minimal entropy is achieved by maximal order, as in crystals and other regular patterns that are anything but complex. Maximal entropy is achieved by minimal order — that is, chaos. If Nature were genuinely out to minimize entropy, She would turn us all into popsicles! Despite there being a distinction, it is certainly not the case, *pace* Wicken (1979), that order and complexity are *opposites*! Chaos is not complex. There is a complex relation between order and complexity which we cannot precisely dissect here, but clearly the “correlation” between them is somewhere between 0 and 1, and that is all that we need for complexity building to be fighting entropy, because order and entropy are indeed opposites.

Maximum entropy is achieved when the probability of system states is uniformly distributed, i.e., every microstate is as probable as every other. This distribution of states is precisely what passive diffusion asymptotes towards for any initial interval on the complexity dimension, as illustrated in Figure 1(c). In other words, the Weak AOC and the passive diffusion argument which potentially explains it immediately reconcile complexity growth with the increase in entropy. This is so immediate that they offer no support whatsoever to the intuition that evolution somehow opposes entropy. If they are the only story to be told of evolution and entropy, then the verdict is rather that that intuition is an illusion. As entropy increases in the evolving system, as states occupied by the system diffuse across the spectrum, the extremal states become ever more extreme. Such is the nature of entropy.

This story is emphatically *not* supporting the increase in complexity of Strong AOC: as the species of an evolutionary system become ever more diffusely spread across the spectrum, the system as a whole becomes ever less organized; it is succumbing to entropy, rather than reversing it!

Passive diffusion could conceivably also support increased complexity of the evolutionary system itself (i.e., Strong AOC). But the only way to work that argument would be to multiply the evolutionary system, so that there are many evolutionary systems that are variations upon one another and which occupy systemic complexity states; over time, as one evolutionary system gives rise to another, randomly more or less complex than itself, some evolutionary systems will inevitably become more complex. If we can realize such a megasystem, then systemic complexity will increase through diffusion, through the operation of a megasystem incorporating many evolutionary subsystems. But this vision is the Gaia hypothesis gone mad! We do not suppose that the universe has implemented this mad vision.

Thus, Gould (1996) is correct to conclude that the passive diffusion argument fails to support any substantial version of the AOC. And, indeed, while supporting Weak AOC it appears to undermine Strong AOC.

5 Transitions, Extinctions and Adaptive Radiation

So what reasons have we to believe in our intuitions, that Strong AOC is right? First of all, we should note that Strong AOC is entirely compatible with thermodynamics. The appearance of a contradiction arises from ignoring the key constraint on the second law: it is about *closed* systems. Evolutionary systems are not closed. Biology gets started on a planet; evolution generates greater systemic complexity for a time; volcanic turbulence or an asteroidal impact

greatly reduces that complexity; evolution through adaptive radiation replaces that complexity; and so on. Until there is a really big extinction event, when evolution no longer operates. Over the long term, entropy will always win. The intuition that entropy is losing is entirely parochial. But, at least according to Strong AOC, that parochial view is spot on in between extinction events: even systemic complexity is then increasing.⁸

We also note that one of the main polemics directed against the consensus view on natural AOC, the ironically titled “Evolution and Complexity: What Everybody Knows” (McShea, 1991), reviews the empirical evidence; however, that evidence exclusively concerns Weak AOC. McShea finds that evidence inconclusive in either direction, with examples both of lineages increasing and of lineages decreasing in complexity; we find that evidence irrelevant to the issue at hand, which is the direction of biosystem complexity change.

There is, however, direct evidence in support of Strong AOC. This evidence is two-fold.

Fact 1 (Major Transitions) *The major transitions in evolution — e.g., from replication via RNA to DNA, from prokaryote to eukaryote, the segregation of the germ line (Dawkins, 1989; Maynard Smith and Szathmáry, 1995) — are not minor amendments to the genetic codes of species, but massive reorganizations of the evolutionary system.*

It seems implausible that passive diffusion alone from pre-existing replicators could produce the major transitions. The preconditions for that argument to work included the ability to lay out the changes on a scalar dimension and offer mutations that will move one up and down that dimension. Minimally, there is no clear prospect of locating the major transitions on such dimensions.

Fact 2 (Mass Extinctions and Adaptive Radiation) *The repeated history of mass extinction events and the subsequent massive rediversification of species strongly indicate that evolution on earth will rapidly occupy any hospitable, unoccupied environments, filling them not with simple copies of existing species and ecosystems, but with complex new ecosystems developed through multiple adaptive radiations.*

It is mass extinction that bars the arrow of complexity. If we measure this history with the (admittedly crude) heuristic stand-in for complexity of the number of species (or genera or families) observed in the fossil record (e.g., Sepkoski et al., 2002; Bambach, 2006), then the complexity growth after mass extinctions is very much like logistic population growth in ecological experiments with fixed energy inputs (Ward, 1995).⁹ A similar kind of process also occurs repeatedly on smaller scales, after forest fires, volcanic eruptions and the like. Plateaus of biodiversity are regularly reached in short order (short on geological time scales). And the degree of biodiversity achieved at those plateaus is highly correlated with known capacity constraints, so, for example, biodiversity levels achieved in tropical jungles (and tropical coral reefs) are systematically higher than those achieved in deserts and savannahs. The point is that immediately after mass extinctions the evolutionary system is massively simplified compared to where it was before. Inevitably, therefore, it has great opportunity to regenerate its lost

⁸There have been parallel discussions of biology and entropy, exemplified by Schrödinger (2003), finding that individual organisms fight entropy, that the developmental process from egg to adult to senescence and death runs against the tide of entropy for the time of one lifespan. Schrödinger focused on the complexity and organization of organisms, but that is quite different from the complexity of organisms involved in the Weak AOC. Schrödinger's organisms were individuals, whereas the “organisms” of the Weak AOC are really *species*. The issues raised are distinct: for Weak AOC, the issue is how diffusion through speciations moves species through the genotypic design space; for Schrödinger, the issue is how ontogenesis produces ordered bodies rather than chaos.

⁹But see Marshall and Jacobs (2009) for a Permian example closer to pure exponential growth.

complexity, building upon the remaining biotic and abiotic resources, since the capacity of the earth (or local ecosystem) to support greater complexity is vastly underutilized. Equally inevitably, according to paleontological evidence (and observations of more local events), the system does grow in complexity as it utilizes that capacity, and it does so with great rapidity.

Also, in an apparent interaction with the major transitions of Fact 1, after the clearing out of the mass extinctions, evolution has often found new ways of utilizing the earth's carrying capacity, for, although it quickly reaches a new plateau of biodiversity, that plateau has often been at a higher level. For example, the coadaptive radiations of insects and angiosperms have resulted in a massive increase in the number of plant species in the Cenozoic (Signor, 1994).¹⁰ The result is that the number of species at the near present (i.e., 30,000 BP) is far greater than during any prior geological era.¹¹

Our observations of the more local events are, to be sure, primarily of repopulations of evacuated areas, rather than the operation of evolution through speciations (e.g., Simberloff and Wilson, 1969). So these events, while analogous to mass extinctions, are not direct evidence in favor of Strong AOC.¹² Nor are they evidence against, since the lack of evolutionary action in such cases is surely due to the lack of an opportunity: it is far too easy and quick for nearby populations to occupy vacant niches for us to see much evolving. The evidence of mass extinctions remains.

Those who would deny Strong AOC simply leave Fact 1 and Fact 2—basic to the story of evolution on earth—unexplained. We shall now proceed to explain Fact 2.

6 Systemic Complexity

The first step is finally to consider directly the problem of defining complexity. The definition of biological complexity already has a considerable history of lack of consensus. We shall treat that history only insofar as the arguments within it are likely to interfere with an appreciation of our underlying approach to understanding biological complexity, which lies within information theory, or the cruder devices we resort to get the practical job done.

6.1 Information-theoretic Complexity

The most versatile and fruitful concept of complexity we know of is information-theoretic complexity, founded on Shannon information theory. And information-theoretic complexity has been widely applied, or widely misapplied, to characterizing biological complexity. Shannon (1948) invented the measure of the amount of information in a message m :

Definition 2 (Shannon Information)

$$I(m) = -\log p(m)$$

This reports the code length required to communicate m when utilizing a communications channel at maximal efficiency; if the log base is 2, this is the bit length of such a message. An efficient code will allocate short code lengths to messages which are common and longer code lengths only to unlikely messages, with the result that the *average* code length is minimized. At

¹⁰One reason for this may well be that mass extinctions select for generalists amongst species; cf. Jablonski (2001).

¹¹An alternative view held by some is that the apparent successive increases in biodiversity is an artifact of biased availability of paleontological data.

¹²However, there certainly are also rapid adaptive radiations in such cases as well (Kocher, 2004).

the extremes, impossible messages have infinite code lengths and certain truths have code lengths of 0 bits — in other words, we don't have to communicate what is already known. The probability-weighted average code length is otherwise known as entropy (Definition 1).

So, the message length of an efficient code is always relative to the probability distribution governing message generation, and that message length is an exact reflection of the amount of information it carries, relative to that distribution. A direct attempt to turn this to measuring biological complexity would be to equate efficient code length with complexity: the more complex an organism is, the longer its description must be! Such attempts have been common (e.g., Hinegardner and Engelberg, 1983; Kampis and Csányi, 1987; Badii and Politi, 1997; McShea, 1991; Bar-Yam, 1997). But these all repeat Wicken's (1979) mistake of identifying complexity with disorder, since the more disordered a system is, the longer its description. Disordered systems may take a long while to describe, but they are not complex in any relevant sense. Biological complexity, or systemic complexity in general, must be measured with a little more sophistication than that. There is, in fact, a highly successful tradition of measuring complexity with information-theoretic tools which does not fall into Wicken's error, namely Minimum Message Length (MML) theory (Wallace, 2005).

MML specifically applies information measures to the job of statistical inference. It does this by dividing messages into two components, one (call it h) which describes the hypothesis under consideration and another (call it e) which describes the sample statistics available (evidence). MML then seeks to minimize the joint message length $I(e, h) = I(e|h) + I(h)$, where $I(h)$ is the Shannon description length of h and $I(e|h)$ is the description length of the evidence assuming that the hypothesis h is true. This is equivalent to maximizing the posterior probability of the hypothesis given the evidence, so MML is a computational approach to Bayesian statistical inference.

To apply MML we need to generate a joint coding scheme for potential evidence and possible hypotheses which is Shannon efficient. If we have a prior probability distribution over hypotheses, we can readily do this using coding theory techniques (e.g., Huffman coding). If we do not have a pre-existing distribution over hypotheses, Wallace advocates codes which treat simple hypotheses as a priori more probable and complex hypotheses as a priori less probable. MML then plays a tradeoff between the complexity of the hypothesis (prior probability) and its fit to the data (likelihood). The simplest possible hypothesis (one that is genuinely null) tells us nothing, so $I(h) = 0$, and, because of that, it requires a complete specification of every detail in the data, so $I(e|h)$ is maximal. The most complex possible hypothesis will have a free parameter for every data point, making $I(h)$ maximal and $I(e|h) = 0$. By minimizing the *joint* length, MML typically finds an intermediate h which has all and only the complexity justified by the evidence available.

We can now see how MML implicitly measures systemic complexity and disorder. It does *not* do so by counting bits over the entire message, but only those dedicated to describing the hypothesis — the *pattern* behind the data, their organizing principle. Totally disordered systems have no organizing principle, so they have zero complexity, and *all* of the joint message is devoted to specifying microstructure. Totally ordered systems, such as crystals and popsicles, usually have extremely simple patterns and little or no variations within the microstructure: both parts of the message are short. Genuinely complex systems are those with complex patterns, large $I(h)$, which normally come with complex microstructures as well (data details), since otherwise the pattern complexity will likely be unjustified.¹³ The failure to distinguish between the uninteresting complexity of microstructure and the interesting complexity of macrostructure

¹³A similar application of information-theoretic complexity to measuring biosystem organization can be found in Chaitin (1979). See also Papertin (1980, 1982).

is behind much of the confusion about biological applications of information-theoretic complexity, such as Wicken's confusion of disorder and complexity (repeated many times, as we noted above), and again behind McShea's (1991; 1996) bizarre suggestion that demolished automobiles and their working counterparts (or living organisms and their dead relatives) are equally complex.

Although MML gives us an *idea* about what a biological complexity measure might be like, it remains unclear how to apply it. In principle, we should have to describe an n-dimensional space within which all morphological and biochemical patterns could be described, along with all possible deviations from those patterns realized in organisms. While MML has been applied to many inferential problems, for example, classification, mixture models, factor analysis and causal modeling, none has been as complex as describing the full functionality of a prokaryotic cell. We shall have to turn to something simpler, such as DNA.

6.2 DNA Complexity

We can directly apply these information-theoretic principles to the so-called "C-value paradox", which will both illustrate the principles in practice and remove much of the paradoxicality of the C-values. DNA codes for ontogeny; it provides the pattern into which the organism grows. Microstructural variations, of course, abound, since environments vary widely. So, it seems organismic complexity has already been directly measured by Nature, and we are almost done: all we need to do is to count up nucleotide bases in the chromosomes. As we have just seen, this has direct support from information-theoretic complexity, and the biological consensus has been: so much the worse for information-theoretic complexity! Why? Because the correlation between apparent behavioral and morphological complexity and genome size (C-value) appears to be nil, with, for example, protozoa having C-values up to 1000 times larger than humans (Gregory, 2004). If we humans have a lower complexity measure than protozoa, then the complexity measure is the problem, not the answer!

There are multiple objections that can be raised against a simple DNA counting method for assessing complexity. One is that we know very well by now that DNA inheritance is not all there is to inheritance. There is a great deal happening with non-genetic inheritances, including cell membranes, mitochondria, the cytoplasm, etc. In addition to that, many species use their wider environment to aid in developing their offspring, what Odling-Smee et al. (2003) call ecological inheritances. There is certainly developmental information carried between generations by both of these methods. Nevertheless, it is also clear that the DNA alone is already sufficient to distinguish between species, so it seems reasonable to infer that the bulk of the information being carried between generations is already in the DNA code. It follows that DNA code length should provide a useful lower bound on species complexity, even while systematically underestimating it, with that underestimation proportional to the complexity of species interaction with the environment.

Another objection is that string length doesn't really matter: what matters is, somehow, *coding* length. We know that much of the chromosome is relatively inactive, with introns not coding for any proteins. So, if we were instead to measure exons, say by counting *genes* rather than bases, we would have a more effective complexity measure. Unfortunately, the reported number of genes per species doesn't support the desired humans-on-top story very well either (Szathmáry et al., 2001). If we take the information-theoretic analysis seriously, we should indeed be looking for *code* lengths, rather than the lengths of this or that string. String lengths are code lengths (in the relevant sense) when and only when the strings are part of an *efficient* coding scheme. It is not at all obvious that there is any very strong selection pressure on

chromosomes for efficiency; indeed, there is likely far more pressure for redundancy in DNA, and thereby ontogenetic reliability, than for efficiency. So, there is no reason to think that chromosomal lengths are informative. What we really need to do is to *compress* DNA strings so that all redundant information is removed and unmeasured; all that would remain would be the real patterns required for ontogeny. Non-functional substrings would then automatically be excluded. In order to maximally compress the DNA we would need to know the real patterns, which we do not. Szathmáry et al. (2001) make some effort in that direction by considering the number and complexity of gene-regulation networks, but their treatment is preliminary and speculative. In fact, researchers have been applying compression algorithms to DNA in an attempt to uncover some of its pattern information (Grumbach and Tahi, 1993; Cao et al., 2007). So far as we know, none of these efforts have been aimed at the C-value paradox; in any case, until they succeed in uncovering and eliminating all the redundancy in DNA, counts based upon them will overestimate DNA complexity. Regardless, even a suboptimal compression technique would give a far better estimate of DNA complexity than has been applied to the C-value question thus far.

6.3 Niche Complexity

There is one final, and we think utterly compelling, objection to DNA length as a complexity measure: it measures, if any thing, the wrong thing. What matters for the Strong AOC — if not for human vanity — is systemic complexity, where the system is the entire biosystem, not an individual organism. A more plausible candidate for system complexity would be a measure of the complexity of the web of niches occupied throughout a biosystem. This is what we propose.

We could nevertheless use DNA complexity to estimate niche complexity, by summing DNA complexities over species within the system. Since DNA codes for both morphology and behavior — i.e., for all functions that the organism fulfills — this should reflect the complexity of the niches which the organism occupies and (by iteration) generates for other organisms. Of course, this will underestimate that complexity, since it does not incorporate complexities resulting from epigenetic and ecological inheritances that have no corresponding encoding in DNA. Regardless, this cumulative measure is likely to be quite informative, assuming a good compression scheme is used.

Niche complexity, of course, should have a more direct measure. This will be a function of both the number of niches and of their interrelatedness. Informally, niches have been described as the roles or loci that organisms occupy (their “ways of life”), having the dual aspect of providing resources (inputs) to the organisms within the niche and in turn providing resources to others outside the niche (outputs or impacts) (Leibold, 1995). If we can identify niches and their input/output relations, we can describe them in a directed graph. Measuring niche complexity can then be achieved by measuring graph complexity, which we know how to do.¹⁴ In case finding and graphing these interrelations is too difficult, we can substitute the heuristic of simply counting niches.

Either way, we need to have a more exact idea of what niches are. The following is the best extant informal definition of *niche* that we know of, one which draws upon and modifies ideas of Hutchinson (1944, 1957) and Odling-Smee et al. (2003):

Definition 3 (Niche) *A niche is “the joint description of the environmental conditions that allow a species to satisfy its minimum requirements so that the birth rate of a local population is*

¹⁴Korb and Nicholson (2004), chapter 8, define a complexity measure for acyclic graphs, which can be readily adapted to cyclic graphs. To be sure, any graph complexity measure would need to be supplemented to account for those niche constraints that do not link niches together.

equal to or greater than its death rate along with the set of per capita effects of that species on these environmental conditions.” (Chase and Leibold, 2003, p. 15)

Niches are thus hyperregions in an n -dimensional space defined using the environmental dimensions along which selection pressures can be identified for a given species, along with any others causally downstream from it. The defining dimensions may be either abiotic or biotic, and, in almost all existing cases, will be both. Note also that this definition incorporates factors describing both the inputs and outputs of niches (the “requirement niches” and “impact niches” of Leibold, 1995). In this we prefer to go slightly farther than Chase and Leibold, interpreting niches as roles in the engineering sense, where functional roles are defined in terms *jointly* of inputs and outputs. So, our interpretation of the definition incorporates the consideration of selection pressures on the occupants of the niche, but also those on the occupants of downstream niches.

Also, contrary to Odling-Smee et al. (2003) and Hutchinson (1978), as well as Chase and Leibold (2003) themselves, we do not understand this definition to be strictly relative to species. In their definitions (or interpretations), species and their niches are logically interdependent: each species has its unique niche and each niche its unique species; an unoccupied niche is a logical impossibility. While we allow for relativization to any species, by observing what selection pressures are actually operating on that species, we do not demand it.

Chase and Leibold (2003) suggest that interdependence of niche and species is the *logical consequence* of the definition above, since it identifies a niche as a property of a species, and you can't have a property of something without having the something. However, this is just wrong if the species referred to is hypothetical or potential rather than real. Counterfactual talk about hypothetical species is not only meaningful, but mandatory. Without it, the attempt to explain the origin of life turns into the classic chicken-and-egg paradox: either we have an infinite regress or a simultaneous miracle. Our answer is simple: the empty abiotic niche came first. Perhaps the hostility of some ecologists for empty niches arises from a failure to note that empty niches are compatible with a weaker form of interdependence: any unoccupied niche may, as soon as it is occupied, become altered by the occupier. Empty niches may well be entirely transitory — without being entirely illusory.¹⁵ Indeed, on our definition the transformation of any actual, but unoccupied, niche into another is the necessary consequence of its occupation, since the new inhabitant will immediately provide new resources to other potential downstream species, in its activities, its detritus and its corpse. As Chase and Leibold (2003) put it, the occupied niche is a result of the interaction between organisms and their environments.¹⁶

In sum, we allow for both purely abiotic niches (“recesses”) and abiotic-biotic niches, and also for niches themselves to evolve, as in ecosystem engineering and the coevolution of species. In any active ecosystem we expect the vast majority of niches will be identifiable only in terms of *both* biotic and abiotic dimensions.

Two properties commonly attributed to niches can hold under our definition:

1. Some species occupy multiple niches simultaneously, e.g., homo sapiens, which is able to adapt to a wide variety of environments.

¹⁵Invasive species provide contemporary examples of empty niches being occupied and transformed, although they are not empty *abiotic* niches.

¹⁶Note that our definition of ‘niche’ incorporates both abiotic and biotic functional roles. Empty niches necessarily are transitory, transforming when occupied, since they thereby acquire biotic functional roles (the downstream per-capita effects), which in the prior abiotic niche are non-existent.

2. Some niches are occupied by multiple species, for example savannah grazing animal species.

Niche construction is where complexity measurement finally makes contact with possible mechanisms for biosystem complexity generation. There is here finally some story to tell about Strong AOC, about how processes of niche construction can possibly yield extremely rapid biosystem complexity growth. We do this with Hypothesis 3 below.

First, however, we note some difficulties involved in measuring the niches we have defined. Since niches are not tied to extant species, and since subsetting niches will yield further valid niches (hyperregions), we must typically have infinitely many niches present in any biosystem — indeed, uncountably many when any defining dimension is continuous (e.g., temperature tolerances). But in considering Strong AOC we are, of course, only interested in *occupied* niches and their graph complexity, not in some unrealizable potential. Thus, we can do what ecologists do, which is to examine the life activities of real species and measure the dimensions of actual selection pressure operating upon them. In practice, this is likely to link niches and species together in the way most ecologists prefer, in a one-to-one mapping. While in theory we think that is not necessary, it may well be practically necessary, and for much of the remainder of the discussion we will accept that linkage.

Another difficulty is that our niche definition is sufficiently vague to accommodate very different degrees of resolution when fixing niche boundaries along one or more dimensions. There will be ambiguities about whether similar hyperregions occupied by two distinct species count as one or as two niches — property (2) above can be made false simply by insisting upon very precise niche boundaries. In order to put the definition into practice it will have to be made more exact, including tolerance parameters defining the precision required for niche boundaries. As we are not sure *how* this should be put into practice, we now review some possible heuristic stand-ins for niche complexity measurement. The point of the review is that the suggested measures will be more or less feasible in different studies under varying circumstances, so having a range of heuristic measures is desirable.

6.4 Heuristic Measures of Niche Complexity

We first consider closer approximations to niche complexity and then successively worse approximations. The closest approximation we suggest is biogeochemical path complexity. Niches are largely, but not entirely, composed of biogeochemical resources consumed by the organism and provided by it to downstream consumers. Measuring energy and matter flow through the biosystem will miss the complexity resulting from environmental tolerances, but should yet provide a useful lower bound on niche complexity. In A-Life simulations it should be relatively straightforward to measure this path complexity, since all paths will be explicitly simulated. To be sure, in most extant A-Life simulations there are no more than a few such paths and they typically are not allowed to be cyclic (e.g., the bodies of organisms do not themselves decompose into their environments), but those limitations are, in fact, at the heart of our prior criticism in Dorin and Korb (2007) and Dorin et al. (2008).

This measure is considerably more accurate than simple food-chain complexity. Some A-Life commentators have noted that measuring food-web complexity might be useful (e.g., Yang, 2002). What these commentators have not noted is that food-webs are a proper subset of biogeochemical path networks, which also include energy and matter exchanges beyond simple consumption, and that biogeochemical paths in turn make a proper subset of paths between niches. Nevertheless, food-web complexity is the next best heuristic and may be chosen should

biogeochemical path complexity be too difficult to measure. In the simpler A-Life simulations, of course, there is no distinction between these two.

As with niche complexity, we have two types of measurement possible: (1) the graph complexity of the biogeochemical path networks (or food webs); (2) a simple count, either of the number of links in the network or the number of its nodes.

The crudest measures of niche complexity we will consider are based on counting the number of species in a biosystem. This is the primary means of estimating biodiversity in use by ecologists today (e.g., Gotelli and Colwell, 2001) and would be the primary means of estimating biodiversity in the past, but for inadequacies in the paleontological evidence, which force us to use the even cruder counts of genera, families and orders (e.g., Benton, 2001; Bambach, 2006).

All of these heuristic measures of complexity apply equally, or misapply equally, regardless of what relation holds between niches and species. Of course, if we accept the numerical identity view, then counting the one is the same as the other; on other views, counting species will misestimate the number of niches. In either case, such simple counts are crude compared to the network complexity measures introduced first.

7 Open-ended Evolution via Niche Construction

We now turn briefly to the relation between niche complexity and the Strong AOC, arguing explicitly that niche complexity in real biology grows exponentially.

Hypothesis 3 (The Arrow of Niche Complexity) *The Arrow of Niche Complexity has an exponential trajectory.*

With complexity interpreted simply as the number of niches, this means that any ecosystem not limited by capacity constraints will robustly tend to produce new niches at an exponential growth rate.¹⁷ We cannot prove this hypothesis, but we can offer three arguments in support.

First, every species creates multiple niches, without exception. The first and most obvious niche is the availability of that type of organism for consumption. Other niches are created by the organism's waste products and its potential for decomposition should it die other than by being eaten directly. In addition to those are typically multiple effects upon the environment via its behavior or simply its presence. Ecosystem engineers rework their environments in predictable ways, ways upon which other organisms can rely for shelter, food, etc. (Jones et al., 1997). In short, any new niche being filled by an organism implies the creation of multiple further niches. This produces a classic "population explosion" of niches.

Our second argument is simply that given the common identification of niches with species, which we have also adopted on practical grounds, Hypothesis 3 follows trivially from the uncontroversial point that biodiversity measured in the number of species increases exponentially up to some saturation level (which, after all, is what led to the open-ended evolution challenge in the first place). To be sure, this argument lacks explanatory power.¹⁸

¹⁷It's worth noting that if we use the more appropriate notion of niche complexity, namely graph complexity, the exponentiality of niche complexity growth only intensifies, since the number of possible graphs is superexponential in the number of nodes.

¹⁸However, if we reopen the relation between niche and species, this argument becomes more interesting, for then we would also have to argue that the number of species and number of niches will track each other even without that identity assumption. Our theoretical rejection of numerical identity was on two grounds: (1) that a strict 1-1 mapping leaves the origin of niches unaccountable; (2) that it is widely accepted that distinct species may occupy the same niche and single species multiple niches. The problem of origins has no implications for the current argument, so we can ignore it here. Regarding the second issue, it's clear that we can tune the limits and tolerances along niche

For our final argument, we have developed a simulation of niche construction which is as simple and general as we can make it, while fitting our account of niches above. Unoccupied niches are cells in a two-dimensional torus, with associated (unutilized) resources. Organisms whose requirements match those resources may occupy these cells, making them occupied niches. Occupied niches are represented as input-output roles, as we have described above: they have varying numbers of input requirements and provide varying numbers of outputs. Over time they replicate, representing the successful production of new species or subspecies by their occupants.¹⁹ Niche mutations allow for changes to either input or output resources, including the addition of new requirements or products and the elimination of others. New niches may arise anywhere in the Moore neighborhood, so long as the new niche's input requirements can be satisfied; if that condition fails, the new niche is still-born. The result of running this simulation is exponential growth in niches under a wide range of initial conditions, in accordance with Hypothesis 3. (For a full description of our simulation and results see Korb and Dorin, 2009.)

These arguments are not of equal weight and significance. Indeed, the second argument, while valid, has no explanatory significance. The first argument seems ineluctable, however, and our simulation reinforces it. That argument implies that the bodies of organisms provide much of the resources fueling the exponential growth of niches, so a precondition for proper niche simulation must be the integration of organisms and environment, sharing biogeochemistry in a *geochemically open organization*, or *goo* for short (see Dorin et al., 2008). The implication for open-ended evolution in artificial simulation is clear: efforts to provoke exponential species growth which don't even try to simulate such niches are hopeless.

8 Conclusion

In the end, it is clear that catching up with the Arrow of Complexity is quite hard: each time you miss you have to redouble your efforts. Nevertheless, we think the effort is worth it, that is, that the challenge of artificially producing open-ended evolution is a worthy one to meet. No attempt thus far has included a serious effort to model niches and their generation, if only because no attempt has gone to the trouble of shaping both environment and organisms from the same basic goo. The complexities and difficulties of doing so while simultaneously modeling an interestingly rich and informative, evolutionarily active biosystem are clear. Indeed, if our argument is even roughly correct, then meeting the goo challenge should suffice to meet the open-ended evolution challenge. These challenges are not equivalent, so there may well be alternative ways to generate open-ended evolution, of course, but no one has yet produced such an alternative. Slogging through the goo seems to us to be our best hope of eventually seeing the fiery release of complexity's arrow unbound.

dimensions so as to merge any two niches, or again so as to place any two distinct regions into separate niches. So, in some sense, anything goes — and we can get any answer we like! But our idea in leaving such matters open was not to endorse something like Feyerabend's (1978) methodological nihilism. Rather, it was to allow ecologists (virtual or not) to be free to accommodate whatever niche taxonomy they may have adopted for other reasons. Reasonable niche taxonomies may vary from the usual 1-1 mapping, but it's not very plausible that the ratio between niches and species would vary by more than some large constant proportion in any single biosystem. (This is particularly so in light of the apparent support niche differences provide biodiversity; Levine and HilleRisLambers, 2009.) But in that case asymptotic complexity theory will show that any constant proportional difference in the number of species and number of niches will vanish in importance relative to an underlying exponential growth rate.

¹⁹Note that in this simulation we make no distinction between niches and their occupying species.

Acknowledgements

We would especially like to acknowledge the contribution of Mark Bedau, both for originally stimulating this research and for extensive discussion of its content. We would also like to acknowledge helpful discussions with Martin Burd, Suzanne Sadedin and Tim Taylor and material assistance from the LSE Centre for Philosophy of Natural and Social Science and the Tilburg Centre for Logic and Philosophy of Science.

References

- Badii, R. and A. Politi (1997). *Complexity: Hierarchical structures and scaling in physics*. Cambridge: Cambridge University.
- Bambach, R. K. (2006). Phanerozoic biodiversity mass extinctions. *Annual Review of Earth and Planetary Sciences* 34, 127–155.
- Bar-Yam, Y. (1997). *Dynamics of complex systems*. Boulder, CO: Westview.
- Bedau, M. (2006). The evolution of complexity. *Symposium on the Making up of Organisms*, Ecole Normale Suprieure.
- Bedau, M., J. McCaskill, N. Packard, S. Rasmussen, C. Adami, D. Green, T. Ikegami, K. Kaneko, and T. Ray (2000). Open problems in artificial life. *Artificial Life* 6, 363–376.
- Bedau, M. A. and N. H. Packard (1991). Measurement of evolutionary activity, teleology, and life. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen (Eds.), *Artificial Life II*, pp. 431–461. Addison-Wesley.
- Bedau, M. A., E. Snyder, and N. H. Packard (1998). A classification of long-term evolutionary dynamics. In C. Adami, R. Belew, H. Kitano, and C. Taylor (Eds.), *Artificial Life VI*, pp. 228–237. MIT Press.
- Benton, M. J. (2001). Biodiversity on land and in the sea. *Geological Journal* 36, 211–230.
- Bronowski, J. (1970). New concepts in the evolution of complexity. *Synthese* 21, 228–246.
- Bullock, S. and M. A. Bedau (2006). Exploring the dynamics of adaptation with evolutionary activity plots. *Artificial Life* 12, 193–197.
- Cao, M. D., T. I. Dix, L. Allison, and C. Mears (2007). A simple statistical algorithm for biological sequence compression. In *IEEE Data Compression Conference*, pp. 43–52.
- Chaitin, G. J. (1979). Toward a mathematical definition of “life”. In R. D. Levine and M. Tribus (Eds.), *The maximum entropy formalism*, pp. 477–498. MIT Press.
- Channon, A. (2006). Unbounded evolutionary dynamics in a system of agents that actively process and transform their environment. *Genetic Programming and Evolvable Machines* 7, 253–281.
- Chase, J. M. and M. A. Leibold (2003). *Ecological niches*. Chicago: University of Chicago.
- Dawkins, R. (1989). The evolution of evolvability. In C. Langton (Ed.), *Artificial Life*, pp. 201–220. Sante Fe: Addison-Wesley.

- Dawkins, R. (1996). *Climbing Mount Improbable*. New York: W.W. Norton.
- Dawkins, R. (1997). Human chauvinism. *Evolution* 51, 1015–1020.
- Dorin, A. and K. B. Korb (2007). Building virtual ecosystems from artificial chemistry. In *European Conference on Artificial Life*, pp. 103–112. Springer.
- Dorin, A., K. B. Korb, and V. Grimm (2008). Artificial-life ecosystems. In *Eleventh international conference on the simulation and synthesis of living systems*, pp. 173–180. MIT Press.
- Feyerabend, P. (1978). *Against method*. Verso.
- Gershenson, C. and R. Lenaerts (2008). Evolution of complexity. *Artificial Life* 14, 241–243.
- Gotelli, N. J. and R. K. Colwell (2001). Quantifying biodiversity: Procedures and pitfalls in the measurement and comparison of species richness. *Ecology Letters* 4, 379–391.
- Gould, S. (1989). *Wonderful life*. New York: Norton.
- Gould, S. (1996). *Full house*. New York: Random House.
- Gould, S. (1997). Self-help for a hedgehog stuck on a molehill. *Evolution* 51, 1020–1023.
- Gregory, T. R. (2004). Macroevolution, hierarchy theory, and the C-value enigma. *Paleobiology* 30, 179–202.
- Grumbach, S. and F. Tahi (1993). Compression of DNA sequences. In *IEEE Data Compression Conference*, pp. 340–350.
- Hinegardner, R. and J. Engelberg (1983). Biological complexity. *Journal of Theoretical Biology* 104, 7–20.
- Hutchinson, G. E. (1944). Limnological studies in Connecticut, VII. A critical examination of the supposed relationship between phytoplankton periodicity and chemical changes in later waters. *Ecology* 25, 3–26.
- Hutchinson, G. E. (1957). Concluding remarks. *Cold Spring Harbor Symposia on Quantitative Biology* 22, 415–427.
- Hutchinson, G. E. (1978). *An introduction to population ecology*. New Haven, CT: Yale University.
- Jablonski, D. (2001). Lessons from the past: Evolutionary impacts of mass extinctions. *Proceedings of the National Academy of Sciences* 98, 5393–5398.
- Jones, C. G., J. Lawton, and M. Shachak (1997). Positive and negative effects of organisms as physical ecosystem engineers. *Ecology* 78, 1946–1957.
- Kampis, G. and V. Csányi (1987). Notes on order and complexity. *Journal of Theoretical Biology* 124, 111–121.
- Kocher, T. D. (2004). Adaptive evolution and explosive speciation: The cichlid fish model. *Nature Reviews Genetics* 5, 288–298.

- Korb, K. B. and A. Dorin (2009). A simulation of niche construction. Technical Report 2009/2, Bayesian Intelligence. <http://www.bayesian-intelligence.com/>.
- Korb, K. B. and A. E. Nicholson (2004). *Bayesian artificial intelligence*. Boca Raton, FL: CRC/Chapman and Hall.
- Leibold, M. A. (1995). The niche concept revisited. *Ecology* 76, 1371–1382.
- Levine, J. and J. HilleRisLambers (2009). The importance of niches for the maintenance of species diversity. *Nature advance online*.
- Lynch, M. (2007). The frailty of adaptive hypotheses for the origins of organismal complexity. *Proceedings of the National Academy of Sciences* 104, 8597–8604.
- Marshall, C. and D. Jacobs (2009). Flourishing after the end-permian mass extinction. *Science* 325, 1079–1080.
- Maynard Smith, J. and E. Szathmáry (1995). *The major transitions in evolution*. New York: W.H. Freeman.
- McMullin, B. (2002). John von Neumann and the evolutionary growth of complexity: Looking backwards, looking forwards. *Artificial Life* 6, 347–361.
- 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., N. Laland, and M. Feldman (2003). *Niche construction*. Princeton: Princeton University.
- Papentin, F. (1980). On order and complexity I. *Journal of Theoretical Biology* 87, 421–456.
- Papentin, F. (1982). On order and complexity II. *Journal of Theoretical Biology* 95, 225–245.
- Ray, T. S. (1990). An approach to the synthesis of life. In *Artificial life II*, pp. 371–408. Addison Wesley.
- Ross, S. (1996). *Stochastic processes*. New York: Wiley & Sons.
- Ruiz-Mirazo, K., J. Umerez, and A. Moreno (2008). Enabling conditions for ‘open-ended evolution’. *Biology and Philosophy* 23, 67–85.
- Schrödinger, E. (1943/2003). *What is life?* Cambridge University.
- Sepkoski, J. J., D. Jablonski, and M. J. Foote (2002). *A compendium of fossil marine animal genera*. Ithaca, N.Y.: Paleontological Research Institution.
- Shannon, C. (1948). A mathematical theory of communication. *Bell System Technical Journal* 27, 379–423, 623–656.
- Signor, P. (1994). Biodiversity in geological time. *American Zoologist* 34, 23–32.

- Simberloff, D. S. and E. O. Wilson (1969). Experimental zoogeography of islands: The colonization of empty islands. *Ecology* 50, 278–96.
- Stanley, S. M. (1973). An explanation for Cope’s rule. *Evolution* 27, 1–26.
- Stout, A. and L. Spector (2005). Validation of evolutionary activity metrics for long-term evolutionary dynamics. In H. G. Beyer and U. M. O’Reilly (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 137–142. ACM Press.
- Szathmáry, E., F. Jordán, and C. Pál (2001). Can genes explain biological complexity? *Science* 292, 1315–1316.
- Wallace, C. S. (2005). *Statistical and inductive inference by Minimum Message Length*. New York: Springer.
- Ward, P. (1995). *The end of evolution: Dinosaurs, mass extinction and biodiversity*. London: Weidenfeld and Nicholson.
- Wicken, J. (1979). The generation of complexity in evolution: A thermodynamic and information-theoretical discussion. *Journal of Theoretical Biology* 77, 349–365.
- Yang, X. (2002). Characterization of multispecies living ecosystems with cellular automata. In R. K. Standish, M. A. Bedau, and H. A. Abbass (Eds.), *Artificial life VIII: Proceedings eighth international conference on artificial life*, pp. 138–141. MIT Press.