

Evolution: The erratic path towards complexity

Nick Barton* and Willem Zuidema*†

*Institute of Animal, Cell and Population Biology

†Language Evolution and Computation Research Unit

University of Edinburgh

June 19, 2003

A central goal of evolutionary biology is to explain the origin of complex organs — the ribosomal machinery that translates the genetic code, the immune system that accurately distinguishes self from non-self, eyes that can resolve precise images, and so on. Although we understand in broad outline how such extraordinary systems can evolve by natural selection, we know very little about the actual steps involved, and can hardly begin to answer general questions about the evolution of complexity. For example, how much time is required for some particular structure to evolve? In a recent paper, Lenski et al. [1] give an intriguing example of how “digital organisms” can evolve. Their work suggests many lines of research, which might shed new light on an old problem.

Complex systems — systems whose function requires many interdependent parts — are vanishingly unlikely to arise purely by chance. Darwin’s explanation of their origin is that natural selection establishes a series of variants, each of which increases fitness. This is an efficient way of sifting through an enormous number of possibilities, provided there is a sequence of ever-increasing fitness that leads to the desired feature. To use Sewall Wright’s metaphor, there must be a path uphill on the “adaptive landscape”.




Figure 1: The adaptive landscape (usually) is a graph of fitness against genotype. Sketched is a hypothetical example, in which alleles at two genes have a continuous range of effects. Both real and digital organisms have, in contrast, a discrete set of possible genotypes involving many more than two genes. Thus, mutations can take them in very many directions. This high dimensionality makes it more likely that there is some path uphill to the “adaptive peak” (See [2])

The crucial issue, then, is to know what variants are available (what can be reached from where) and what is the fitness of these variants. Is there a route by which fitness can keep increasing? Population genetics is not much help here. Given the geometry defined by mutation and recombination, and given the fitnesses, we can work out how a population will change, simply by following the proportion of different types through time. But understanding

how complex features evolve requires plausible models for the geometry of the adaptive landscape, which population genetics by itself does not provide.

“Artificial Life” — the study of life as it could be — provides a variety of such models. For instance, Thomas Ray [3] developed a model, called “Tierra”, where digital creatures are little computer programs that copy themselves and compete with each other for memory and processing time. Fitness here — just as in the real world — is defined very indirectly by the rate of self-replication of the creatures relative to others. Ray’s creatures evolved strategies to hinder competitors and even to parasitize other creatures. Karl Sims [4] created a simulated physical world in which “digital creatures” successfully evolve both their bodies and brains in order to beat other creatures in a variety of tasks such as swimming, walking and jumping. Lipson and Pollack [5], in a recent follow-up study, actually made such walking creatures as little robots and showed that the evolved locomotion strategies work even in the real world. Fitness in these models is defined implicitly by the complex relation between brain and body architecture and the resulting way of moving.

In Lenski et al.’s recent study, the creatures consist of a string of instructions, each instruction being chosen from 26 possibilities. Like Ray’s creatures, the instructions must implement self-replication in order for the creature to have offspring. But like Simm’s creatures, they are also rewarded for performing a specific task: they can replicate faster by manipulating information from the environment. Each organism receives two random 32 bit strings as inputs, and is rewarded if it produces an output string that matches one of 9 possible logical operations. For example, the logical operation NAND (“not and”) returns a 0 in the output string only if the corresponding digits in the input strings are both 1, and a 1 in all other cases. One of the 26 possible instructions in a creature’s “genome” is a logic operation (NAND), whilst the others perform various manipulations: copying, input/output, etc. Composite logic operations are valued according to the number of elementary NAND operations needed to perform them. The most valuable is EQU (“equal”), which returns a 1 only if both input bits are the same (this requires 5 NAND operations) as well as other operations with more intermediate results between registers. A hand-written program required 19 operations to achieve EQU; a digital organism needs additional code for replication.

Initially, 3600 identical organisms were set up, each with 15 instructions that allowed replication, plus 35 dummy instructions. In each replication, point mutations occurred at a rate 0.0025 per instruction, and single-instruction insertions or deletions occurred at a rate of 0.056 per genome. In one run, EQU evolved after 111 steps. (A “step” is counted whenever offspring differed from parent along the successful lineage; in most cases steps corresponded to single mutations, but 8 steps involved two or three mutations). Over a further 233 steps, the ability to perform additional logic operations evolved, and so fitness increased further. The way in which these organisms evolved was broadly as one would expect. In particular, the evolution of EQU depends on there being fit steps that lead up to it, as allowed by the reward system shown in Table 1.

Function name	Logic operation	Reward
NOT	$\neg A$	2
NAND	$\neg(A \text{ and } B)$	2
AND	$A \text{ and } B$	4
OR_N	$(A \text{ or } \neg B)$	4
OR	$(A \text{ or } B)$	8
AND_N	$(A \text{ and } \neg B)$	8
NOR	$\neg A \text{ and } \neg B$	16
XOR	$(A \text{ and } \neg B) \text{ or } (\neg A \text{ and } B)$	16
EQU	$(A \text{ and } B) \text{ or } (\neg A \text{ and } \neg B)$	32

Table 1: Rewards for performing logical operations. The symbol \neg denotes negation (“not”). Logic operations are performed digit by digit on one or two input strings. Thus, when applied to the input strings “110” and “011”, the operation AND would yield “010”.

Lenski et al. experimented with the computer model in much the same way that geneticists experiment with model organisms, by changing the fitness regimes and by knocking out instructions on the evolved genomes one at a time to test their effect on fitness. They can also do something that geneticists usually cannot: trace back the evolutionary history of the genome that first produced EQU. From the results from this study, Lenski et al. emphasise one feature in particular: often, deleterious changes are established along the path to evolution of EQU. From a population genetics point of view this result is less surprising than it may seem at first sight. One expects some deleterious mutations to

be picked up by random drift in a population of only 3600 organisms. Moreover, these digital organisms are asexual, so that a deleterious mutation can be established if it occurs together with a favourable mutation (hitch-hiking, [6, 7]), or if a new mutation occurs that produces a fit genotype when *combined* with the initially deleterious mutation. In the example analysed by Lenski et al., most of the deleterious mutations along the lineage leading to EQU only reduced fitness slightly, by less than 3%. However, two reduced fitness by more than 50%, and were only rescued by mutations which occurred immediately afterwards — in one case, by the mutation which first produced EQU. Moreover, that evolution of EQU *required* the previous mutation, which initially greatly reduced fitness. This pattern, of strong epistatic interaction, was seen in the final stages of 3 of the 23 replicates in which EQU evolved.

So, in these simulations adaptation frequently depends on the occurrence of double mutations, either in the same generation, or in close succession. Suppose that a particular deleterious mutation arises at rate μ_1 and reduces fitness by s . It is expected to persist for $\sim 1/s$ generations [8] during which time mutations at another locus occur at rate μ_2 . If both occur together, they confer a strong advantage, and are picked up by selection. So, we expect a rate of accumulating these interacting pairs of $\sim \mu_1\mu_2/s$, compared with $\sim \mu_1$ for single favourable mutations. The observation that interacting pairs do get established quite frequently tells us something about the relative abundance of paths involving single mutations versus double mutations: possibly, once all single-mutation steps have been explored, the population must wait until the rarer doublets arise.

In Lenski et al.'s artificial organisms, the mutation rate per site is quite high (0.0025) so that favourable pairs can be picked up by selection at an appreciable rate; this would be unlikely in most real organisms, because there, mutation rates at each locus are low. There are, however, some biological examples in which double mutations contribute to adaptation — the first deleterious, the second favourable in combination. In general terms, Manfred Eigen has argued that evolving populations of RNA molecules form a “quasi-species” [9], with high diversity maintained by predominantly deleterious mutation away from a wild-type sequence that is itself vanishingly rare. This diversity allows the population to explore a larger fraction of sequence space. More specifically, the secondary structure of rRNA molecules can be determined through the pattern of covariation of substitutions: if one base changes, its partner changes soon after in order to maintain base pairing. Here again, the first change occurs by chance, in opposition to selection, and is compensated by the second [10]. Lenski et al. do not explore the applicability of their model to such issues.

Artificial Life models such as Lenski et al.'s are perhaps interesting in themselves, but as biologists we are concerned here with the question what Artificial Life can tell us about real organisms. The difficulty in answering that question is that much work in this field is rather isolated from traditional evolutionary biology. Well-established theories and methods from population genetics and game theory are too often ignored — and Lenski et al., although they explore the evolutionary dynamics in quite detail, are no exception. There are, however, ways in which Artificial Life can benefit from evolutionary theory, and vice versa. Can we understand exactly how complexity evolves in these artificial models? Can we find general rules which describe the process? For example, could we predict how long it is likely to take for a function such as EQU to evolve, given mutation rates and fitnesses? Here, there are population genetics principles which are helpful: the relative rates of single vs. double mutations that we discussed, ideas about “hitch hiking” [6] and Haldane's “cost of selection” [11], and so on. Since the entire fossil history of digital organisms is preserved in the computer, it really should be possible to understand their evolution in quantitative terms.

But conversely, there are also potential benefits for evolutionary biology. In population genetics and evolutionary game theory we design models to study the success and failure of a predefined set of traits or strategies in the struggle for life. But what are the possible traits? And how well do they work out in *particular* environments with *particular* competitors? These questions are ignored in traditional models — they come in as parameters to be provided by developmental biology and ecology. For understanding the evolution of complex traits this is not satisfactory, because these parameters are themselves shaped by evolution. Evolutionary processes constantly shift the targets of evolutionary optimization, create spatial patterns, turn competitors into mutualists and create new levels of selection. Artificial Life models of such phenomena (e.g. [12, 13, 14]) promise to be useful for developing the concepts and techniques to deal with that challenge, but only if they are combined with the insights from almost a century of population genetics.

References

- [1] R.E. Lenski, C. Ofria, R.T. Pennock, and C. Adami. The evolutionary origin of complex features. *Nature*, 8(423):139–144, May 2003.
- [2] W.B. Provine. *Sewall Wright and evolutionary biology*. University of Chicago, Chicago, IL, 1986.

- [3] T.S. Ray. An approach to the synthesis of life. In C.G. Langton, C. Taylor, J.D. Farmer, , and S. Rasmussen, editors, *Artificial Life II*, pages 371–408. Addison-Wesley, 1992.
- [4] Karl Sims. Evolving virtual creatures. In *Computer Graphics (SIGGRAPH '94 Proceedings)*, pages 15–22. ACM Press, New York, NY, 1994.
- [5] Hod Lipson and Jordan B. Pollack. Automatic design and manufacture of artificial lifeforms. *Nature*, 406:974–978, 2000.
- [6] W.G. Hill and A. Robertson. The effect of linkage on limits to artificial selection. *Genetical Research*, 8(269-294), 1966.
- [7] J. Maynard Smith and J. Haigh. The hitch-hiking effect of a favourable gene. *Genet.Res.*, 23:23–35, 1974.
- [8] J. B. S. Haldane. *The causes of evolution*. Longmans, New York, 1932.
- [9] Manfred Eigen. Self-organization of matter and the evolution of biological macro-molecules. *Naturwissenschaften*, 58:465–523, 1971.
- [10] W. S. Stephan and D. A. Kirby. Rna folding in drosophila shows a distance effect for compensatory mutations. *Genetics*, 135:97–103, 1993.
- [11] J.B.S. Haldane. The cost of natural selection. *Journal of Genetics*, 55:511–524, 1957.
- [12] W. Daniel Hillis. Coevolving parasites improve simulated evolution as an optimization procedure. *Physica D*, 42:228–234, 1991.
- [13] Maarten C. Boerlijst and Paulien Hogeweg. Spiral wave structure in pre-biotic evolution: hypercycles stable against parasites. *Physica D*, 48:17–28, 1991.
- [14] R.A. Watson and J.B. Pollack. Symbiotic composition and evolvability. In Jozef Kelemen and Petr Sosik, editors, *Advances in Artificial Life*, volume 2159 of *Lecture Notes in Computer Science*, pages 480–490. Springer, Berlin, 2001.