

Punctuation Equilibrium and Optimization: An A-life Model

Ravi Jonnal¹

Cognitive Science Program
Department of History and Philosophy of Science
Indiana University
Bloomington, IN 47405, USA
rjonnal@indiana.edu

Anthony Chemero²

Scientific and Philosophical Studies of Mind Program
Franklin and Marshall College
Lancaster, PA 17604, USA
a_chemero@acad.fandm.edu

Abstract

In this paper, we discuss an A-Life experiment in which we evolve artificial neural network controllers for an animat. We compare fitness scores of for trials in which the neural networks are evolved with and without punctuated equilibria. Finally, we discuss the way our findings relate to current debates in evolutionary biology.

1. Introduction

Since the spring of 1999, we have been attempting to use the methods of artificial life to study modularity (Fodor 1983), roughly the idea that the mind is not a single general-purpose device, but rather is made up of multiple special-purpose components, called modules, which interact to produce our intelligent behavior. To this end, we designed an artificial creature that lives in a simulated environment. The creature's behavior is controlled by an artificial neural network, a small simulated brain. We have been experimenting with different evolutionary algorithms, attempting to transform the initially randomly-connected neural network into a suitable controller of the creature's behavior.

One of the main convictions guiding our research was that most evolutionary algorithms were not sufficiently beholden to what is known about actual biological evolution. (An exception to this is the evolutionary robotics group at the University of Sussex. See Harvey *et al.* 1997.) Thus it was important to us that our evolutionary algorithm be more realistic than others that have been used so far: we have taken into account that animals change aspects of their environment and, most importantly for present purposes, we have allowed for the possibility of *punctuated equilibria*, sudden changes in

the in rate of evolutionary change itself.¹ We now think that our findings concerning punctuated equilibrium overshadow those concerning the evolution of modularity. It is these findings, those concerning punctuated equilibrium, that we will discuss here.

Ever since Eldridge and Gould (1972) introduced punctuated equilibrium in 1972, the exact nature of its relationship to adaptationism and the rest of evolutionary biology has been a source of controversy (see for example Gould and Eldridge 1993; Dennett 1995). Some (e.g. Gould 1980) have suggested that punctuated equilibrium is inconsistent with adaptationism; Others (e.g. Dawkins 1986) have suggested that careful adaptationists can take punctuated equilibrium in stride. We will not go into the details of the arguments and counter-arguments here, fascinating though they may be (Eldridge 1985 provides a good, though biased, account). Instead, we will say just enough about what the controversy is supposed to be about so that the evidence from our simulations can be seen to undermine it.

Adaptationism is a methodology in evolutionary biology whereby one treats an organism as a set of traits, each of which is selected for because it is the optimal structure available to selection for the fulfillment of its function (see Dennett 1995; Chemero 1998). According to adaptationism, natural selection is the main force behind evolutionary change. Natural selection, it is claimed, causes gradual change to species so that:

- (i) the members of the species become highly adapted to their environments; and
- (ii) if there is reproductive isolation, gradual adaptation to varying environments (or niches

¹ Bak and Sneppen 1993 describe a computer model of punctuated equilibrium. Our work differs from theirs in that for us punctuated equilibrium is just part of an overall evolutionary model. The entire purpose of the Bak and Sneppen simulation is to model punctuated equilibrium.

within the same environment) can lead to speciation.

According to the proponents of punctuated equilibrium (Eldridge and Gould 1972; Gould and Eldridge 1993), for the majority of time species are in evolutionary stasis, with little or no change occurring and hence little or no increase in adaptation to their environments. Occasionally, often due to some environmental catastrophe, there will be punctuations, periods of rapid evolutionary change during which speciations occur. Evolution, according to punctuationalist, is not gradual adaptation to the environment, but proceeds in occasional fits and starts, without necessarily increasing fit between animal and environment. As Michael Ruse (1992) has put it, it is hard to imagine how the sort of discontinuous and very rapid change that occurs in punctuations could "stay in adaptive focus". Thus, punctuated equilibrium seems to amount to a denial of both the claims of adaptationists outlined above. On (i) punctuationalists claim that, except when speciation occurs, species are in stasis and do not become increasingly adapted to their environments; on (ii), punctuationalists claim that gradual natural selection alone is insufficient for speciation, which requires a punctuation event. This apparent conflict led Gould to say (infamously) that "the new synthesis is effectively dead" in 1980.

In this paper, we describe experiments in artificial life in which a neural network is artificially evolved to control a virtual creature. With the evolutionary algorithm employed in the artificial evolution, it is possible to simulate punctuated equilibrium. The aim of the experiments we describe is to compare the fitness in animats whose evolution includes punctuations with the fitness of animats whose evolution does not. We will argue that the results of these experiments call the apparent opposition between punctuated equilibrium and adaptationism into question.

2. Methods

2.1 *The animat and environment*

The model used here was designed for two purposes: to investigate the effects of punctuated equilibrium on artificial evolution and to investigate the conditions under which modularity evolves (or fails to evolve) in an artificial system. Only the first is discussed here. The model consists of a simulated organism (Bukowski, hereafter) and its behavior in a

simple, but changing, environment. The environment consists of a two-dimensional grid of $n \times m$ squares, each of which may be occupied by an obstacle, a resource, Bukowski, or a waste product. The grid is bounded on all four sides by obstacles.

At the start of each epoch (an epoch consists of 200 time steps), Bukowski is given a random position and direction in the grid, and a fixed number of resources and obstacles are randomly distributed throughout the grid. Over the course of the epoch, Bukowski moves around the environment, consuming resources, colliding with obstacles, leaving a messy trail of waste products, and occasionally moving through one of the latter. The frequency of each of the previous events is recorded. Whenever a resource is used, another is placed at a random position in the environment. The waste products remain in the environment for an average of 10 time steps, after which they are removed, regardless of whether Bukowski moves over them or not.

In the interest of computational tractability, we decided to simulate some aspects in very general and sketchy ways. Bukowski consumes an object whenever it is in the same position as the object. It may consume resources or waste products. Similarly with obstacles, Bukowski collides with an obstacle whenever it attempts to move into the obstacle's position; it of course cannot consume an obstacle.

2.2 *The controller*

Bukowski is controlled by a three-layer feed-forward artificial neural network (ANN). It has seven input nodes, which correspond to each of the organisms seven simulated sensors: three obstacle sensors, two resource sensors, and two waste sensors. The ANN has twenty nodes in the second, hidden, layer, none of which are connected laterally. It has two output nodes, which control the organism's movement. The input and hidden layers are fully connected (i.e. each input node is connected to each hidden node), as are the hidden and output layers. There are a total of 29 nodes and 180 connections. The nodes have activation values in the range $[0, 1]$ and the connections have weights in the range $[-1, 1]$, where positive values signify excitation and negative ones inhibition.

At the start of each time step, the simulator uses the position of Bukowski, the configuration of the environment, and certain information about the sensors to determine which objects in the environment affect which sensors. It then computes

an activation value for the sensor based upon how many objects are affecting it and from what distances, and then adds a sensor/environmental noise term which changes the computed activation by a random number in the range [-0.1, 0.1].

The activation of each hidden and output node is determined by the activations of input and hidden nodes, respectively, and input to hidden weights and hidden to output weights, respectively, by applying the logistic function to sums of relevant products:

$$a_j = \frac{1}{1 + e^{-\sum_{i=0}^n a_i w_{ij}}}$$

where a_j is the activation of node j , w_{ij} is the connection weight between nodes i and j , and n is the number of nodes in the relevant layer.

The activations of the two output nodes are used to compute Bukowski's new position and orientation, using a rough simulation of two stepper motors. The sensors are updated once more, and so on.

2.3 The evolutionary algorithm

Bukowski's genome consists of an array of 180 floating point numbers in the range [-1, 1], which are the connection weights of the ANN.

Initially, a population of animats is generated, each with a randomly generated genome. Each animat is placed in different randomly generated environments, for one epoch each, and the numbers of resource-consumptions, waste-consumptions, collisions, and movements are recorded, and the animat's overall fitness is computed with a fitness function (see below). When all animats in the population have been similarly 'tested', the genomes of the 10% with the highest overall fitnesses are used to generate a new population of animats, each one being responsible for a fixed number of offspring.

The offspring are created by copying the parental genome and allowing for mutations. Two sorts of mutations are simulated: small mutations, which alter a single value randomly by a number in the range [-0.1, 0.1], and large mutations which randomly reassign the weight a value in the range [-1.0, 1.0]. To allow for the possibility of entirely severed connections, any new weight between -.08 and .08 is reassigned a value of 0.

The overall rate of mutations μ determines both the rate of large and small mutations. After trying several ways of computing the rates, we

found that using $\mu - \mu^2$ for the rate of small mutations and μ^2 for the rate of large mutations produced interesting results. With some informal experimentation, we determined that setting μ to 0.04 resulted in larger overall fitness gains than $\mu = 0.01$, $\mu = 0.02$, or $\mu = 0.06$, and so chose $\mu = 0.04$ to be our base mutation rate.

So, the next generation is populated with mutated parental genotypes, and same steps are performed on it.

2.4 The fitness function

Here it is:

$$F = 500 \frac{r}{m} - 500 \frac{w}{t} - \frac{c}{t}$$

where r , w , c , m , and t represent the numbers of resource-consumptions, waste-consumptions, collisions, movements, and time steps, respectively, over the course of an individual animat's lifespan. The motivation for using movements in the first term rather than time steps is the fact that in real biological systems, movements are costly.

2.5 Punctuated mutation rates

In some trials, instead of keeping μ constant for the entire trial, we introduced a probability p that μ increases by some factor m over the course of a trial, such that for an individual generation, there is probability p that the mutation rate is set to $m\mu$.

We ran punctuated trials with many different values of both p and m , but will only discuss trials where $p=.02$ and $m=2$, such that the maximum mutation rate is 2μ .

2.6 The experiment

We ran the trials in pairs, pairing one nonpunctuated trial with one punctuated trial. In each trial, a population of 100 animats was evolved for 200 generations. Each individual animat is put into five different randomly generated environments for one epoch each. For the nonpunctuated trials, the mutation rate μ was set to .04, and for the punctuated trials the base mutation rate was the same, while the punctuation frequency p was set to .02 and the punctuation magnitude m was set to 2.

3. Results

In all cases, the artificial evolution improved the Bukowski's fitness. This was evident in two ways: first, in every trial the fitness scores improved during evolution; second, it was there was vast improvement in the creature's overall behavior. Initially, with randomly generated networks, the animats were unable to avoid obstacles and waste or find food. Indeed, some were unable to move at all. After evolution, however, the creatures do seek out food sources and rarely bump into obstacles. This result is unsurprising.

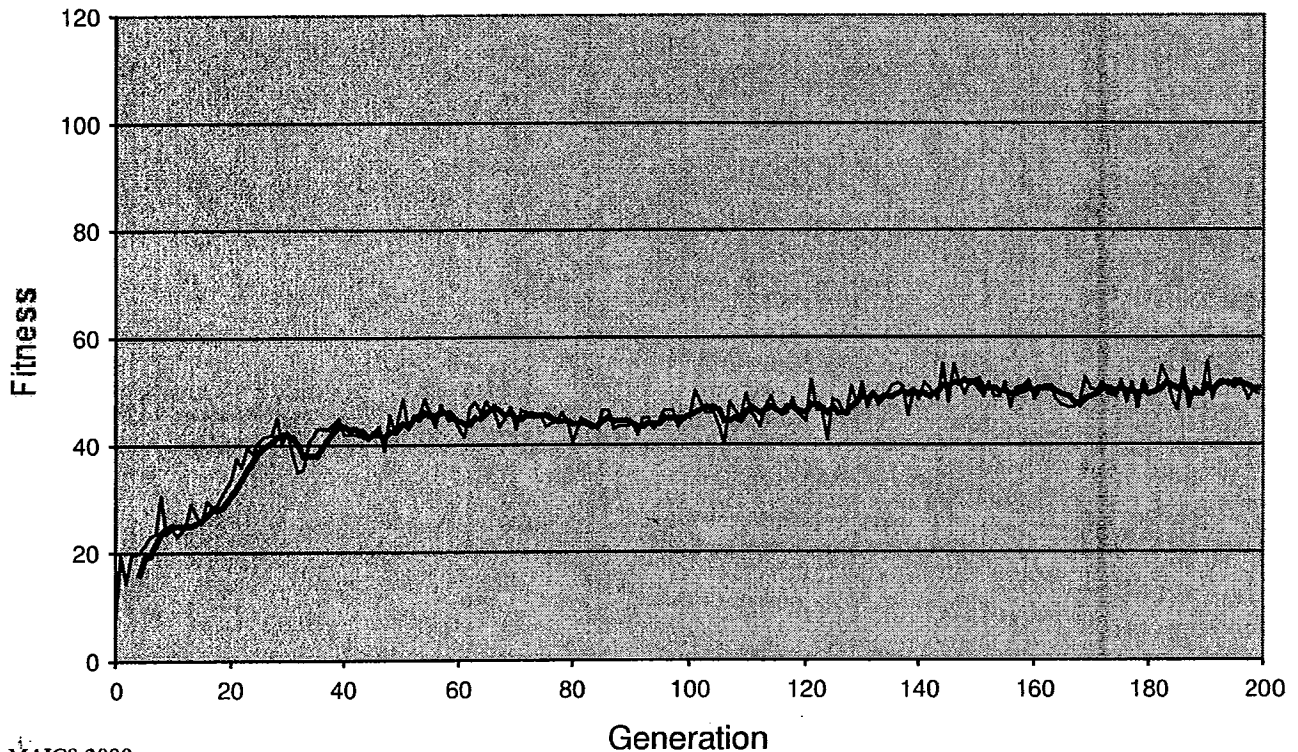
A surprising result is that there was a notable difference between the degree to which fitness improved in the punctuated and non-punctuated trials. In all but one case, the trials that included occasional punctuations had final fitness scores that were better than the scores of the trials that had no punctuations. Also, in all but one case the change in fitness during evolution was greater in the punctuated cases. In one pair of trials (trials 1-control and 1-punctuated), in fact, the creature from the punctuated trial had a final fitness score that was more than twice as high as the control, despite the fact that its initial fitness score was slightly lower. (See Table 1.)

Trial	Initial Fitness	Final Fitness	Fitness Change
1-control	7.64	49.55	41.91
1-punctuated	6.32	109.58	103.26
2-control	6.47	142.81	136.34
2-punctuated	7.74	161.87	154.13
3-control	4.79	46.68	41.89
3-punctuated	6.43	52.18	45.75
4-control	4.67	40.14	35.47
4-punctuated	10.41	53.26	42.85
5-control	3.28	45.53	42.25
5-punctuated	5.42	43.07	37.65

Table 1

The mean final fitness and mean change in fitness were considerably higher in the punctuated trials. This is summarized in Table 2.

Figure 1: Trial 1 Contro



Trial Type	Mean Initial Fitness	Mean Final Fitness	Mean Fitness Change
control	5.37	64.94	59.57
punctuated	7.26	83.99	76.73

Table 2


It is worthwhile to look more closely at the way the fitness changes over time in the actual trials. Because the differences between the control and punctuated trials are most striking in the first pair, we will focus our attention there. Trial pairs 2-4 differ from pair 1 in degree only. Trial pair 5, in which the final fitness and overall fitness gains are greater in the non-punctuated trial, will be discussed below.

In trial 1-control (see Figure 1), there is initially a rather steep gain in fitness. From generation 0, in which the brain has been randomly generated, to generation 30, the fitness increases 34.70 points, from 7.64 to 42.34. Then, after a slight decrease (4.54 points) caused by an unfortunate mutation, the fitness changes are less dramatic.

From generation 30 to generation 199, the overall fitness ranges from a low of 34.84 to a high of 55.26, reaching a final value 49.55. As is evident from the rolling average depicted in figure 1, there is an overall upward trend for these generations, but the slope is rather modest. This is exactly what an adaptationist such as Dennett or Dawkins would expect in non-artificial evolution: as time passes, the evolutionary algorithm yields a slow, but steady, increase in adaptive fit between creature and environment.

Compare this to trial 1-punctuated, in which in 2% of the generations have doubled mutation rates. In this trial, as in 1-control, there is an initially steep gain in fitness (see Figure 2). But unlike in the control trial, this overall trend continues, despite occasional large decreases in fitness. A notable example of this is the punctuation that occurs at generation 166, where fitness suddenly drops 44.18 points from 96.84 to 52.66, a decrease of more than 45% in one generation. Sudden drastic changes such as this one are like the speciation events that punctuationists such as Eldridge and Gould (1972) discuss. And just as they suggest, there are two different sorts of change that occur: normal, slow increases in fitness and sudden, more drastic changes like that which occurs at generation 166. But notice that after the punctuation occurs, the fitness gains are dramatic, and the creature that results is significantly more fit. In this case, the final fitness was 109.58.

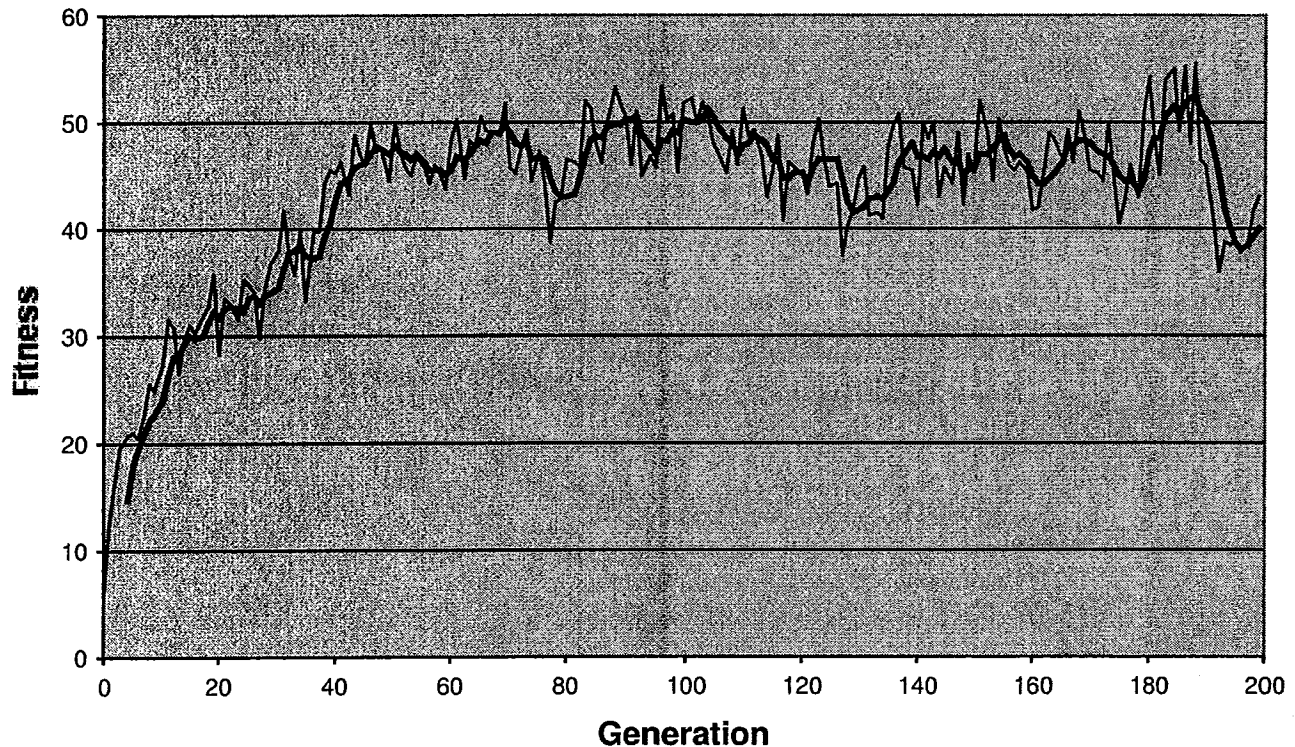
Figure 2: Trial 1 Punctuated



Adding punctuations to the trials had similar, though less dramatic, effects in trial pairs 2-4. In trial 5, however, there was a sudden fitness decrease of more than 15% at the 188th generation of the punctuated group (from 55.40 to 46.51). The remaining 12 generations were not sufficient to make up for the losses (See figure 3). Because of this, the final fitness was lower in the punctuated trial than in the control trial in this pair. But note that the maximum fitness in the punctuated trial (55.40), which was the result, in part, of earlier punctuations, was higher than the maximum fitness of the non-punctuated trial (51.34) in this pair.

4. Discussion

Figure 3: Trial 5 Punctuated



Our results suggest that in artificially evolving systems, punctuating periods of relatively neutral evolution with brief increases in mutation rate can result in better overall fitness gains. That is, punctuated equilibrium can be used as an optimizing tool in evolutionary algorithms research and application. Although this is just speculation, we think that punctuations helped improve overall fitness by "bumping" the evolving system out of

local maxima, allowing to get to more global maxima. Whether or not this is the reason that punctuations tend to improve optimization, our research clearly indicates that they do.

This apparently purely practical result has implications for theoretical questions in evolutionary biology.

In particular, it raises doubts about the supposed opposition between adaptationism and punctuated equilibrium. If, as our results suggest, it is possible for punctuations in evolving systems to produce net increases in fitness, adaptationists might begin to see punctuated equilibrium not as an obstacle to their methodology, but rather as a way in which the adaptations they seek might occur.

It is often thought that punctuated equilibrium implies that not all (perhaps not many) traits of organisms are adaptations to their environments. Just as gradual adaptive change increases the degree of fit between the animal and its environment, so too, in the long run, do changes caused by punctuation. In fact, in the long run, punctuation along with gradual change may result in better fit than gradual change alone. Our results indicate that far from being a hindrance to

adaptation, punctuations can in fact be one of its motors. If our findings in this limited artificial context have a real, biological correlate, adaptationists would no longer need to make apologies or excuses for the apparent existence of punctuations.

These findings are surprising, and we ought to admit that we find them somewhat troubling. Our initial motivation for modeling punctuated equilibrium was what we take to be the overly adaptationist leanings of much of the research in artificial life and evolutionary robotics.² And now our findings might be interpreted to mean that the adaptationists were right all along, and for more reasons, it turns out, than even they realized.

One way to question these implications would be to suggest that our model of punctuated equilibrium is inappropriate in one way or another. It might be argued, for instance, that what causes punctuations in natural biological systems is environmental upheaval, and not sudden and temporary changes in the mutation rate. Maybe. But the overall effect is very similar. In both environment driven and mutation rate driven punctuated equilibrium, there are brief periods of poor fit between animal and environment, followed by longer periods of gradual adaptation and also neutral evolution. Were the poor fit in our model caused by simulated environmental catastrophe, would the results have been different? We have doubts that they would, but a more confident answer will require more research. (We are currently conducting these experiments.)

It might alternatively be argued that what causes punctuations is the founder effect, as suggested by (Gould and Eldridge 1993). When a small subpopulation is geographically isolated from the original population, the subpopulation does not represent the allele frequencies present in the original population, and there are new or different selective pressures on the subpopulation, rapid evolutionary change may occur. It is true that the punctuations in our model are not caused by isolation of subpopulations. However, there is an interesting similarity between brief surges in mutation and what occurs in the founder effect. In both cases, there have been random changes in the genetic makeup of the population. In the founder effect this is due to the small sample size; in our model it is due to the increase in the mutation rate.

This potential objection, like the one before it, can only be answered definitively by more simulation. In the meantime we might wonder whether the apparent conflict between adaptationism and punctuated equilibrium is merely apparent. We might also wonder whether the protracted and often bitter argument over this matter has been just so much hot air.

Acknowledgements

Jonnal would like to thank the Indiana University Cognitive Science Program for funding this research. Chemero would like to thank the Franklin and Marshall College Committee on Grants.

Works Cited

- Chemero, A. (1998). Teleosemantics and the critiques of adaptationism. *Evolution and Cognition*, 4: 136-144.
- Bak, P. and Sneppen, K. (1993). Punctuated equilibrium and critically in a simple model of evolution. *Phys.Rev. Letters*, 71: 4083-408
- Dawkins, R. (1986). *The Blind Watchmaker*. New York: Norton.
- Dennett, D. (1995) *Darwin's Dangerous Idea*. New York: Simon and Schuster.
- Eldridge, N. (1985). *Timelines*. New York: Simon and Schuster.
- Eldridge, N. and Gould, S. (1972). Punctuated equilibria: An alternative to phyletic gradualism. In Schopf, T. (ed.) *Models in Paleontology*. San Francisco: Freeman Cooper.
- Fodor, J. A. (1983) *The Modularity of Mind*. Cambridge: MIT Press.
- Gould, S. (1980). Is a new and general theory of evolution emerging? *Paleobiology*, 6: 119-130.
- Gould, S. and Eldridge, N. (1993). Punctuated equilibrium comes of age. *Nature*, 366: 223-227.

² Again, excepting the University of Sussex group.

I. Harvey, P. Husbands, D. Cliff, A. Thompson, N. Jakobi. (1997). Evolutionary robotics: the Sussex approach. *Robotics and Autonomous Systems*. 20: 205--224.

Ruse, M. (1992). Darwinism. In Keller, E. and Lloyd, E. (eds.), *Keywords in Evolutionary Biology*. Cambridge: Harvard University Press.