

Artificial Death¹

Peter M. Todd

The Rowland Institute for Science
100 Edwin H. Land Boulevard
Cambridge, MA 02142
ptodd@spo.rowland.org

Abstract

We have developed an open-ended system in which to study the evolution of behavior in a population of simulated creatures. The population grows when individual creatures actively reproduce; population size is kept in check by the death of creatures that run out of energy. This allows new creatures, with new behaviors, to have access to the environmental resources they will need to survive, so that constant turnover of individuals and consequent evolution can take place in the population. In many instances, though, super-individuals can evolve that choose to opt out of the energy-depleting reproduction process, becoming for all purposes immortal and thereby stalling out the course of evolution. To solve this problem of immortality, new forms of death and senescence must be added. Differences in the simple rules used for reaping functions can have widely varying effects on the evolved behaviors of individuals and the global behavior of the population as a whole, and carry with them implications for the evolution of death in natural systems.

Introduction

Death can be thought of as nature's way of giving somebody else a chance. As Malthus (1803/1992) observed nearly two centuries ago, a reproducing population unchecked by mortal influences will grow geometrically to fill its environment and exhaust its resources, so that no new growth (introduction of new individuals) is possible. Waddington (1982) has summarized this notion slightly differently: "A species whose individuals were immortal would exhaust its possibilities for future evolution as soon as its numbers saturated all the ecological niches suitable for its way of life" (p. 649). Death reduces the number of organisms competing for the same resources, and allows thus the introduction of new individuals into a population with some hope of survival. These new individuals, with their potential differences from past individuals, can alone provide the raw variance on which ongoing evolutionary processes depend.

¹Work in progress.

As a result – at least in non-expanding populations – death is essential to the ongoing course of evolution.

This conclusion, while relatively obvious, is usually overlooked or understated in simulations of evolution. In the most commonly used genetic algorithm methods (Holland, 1975/1992; Goldberg, 1989), for instance, a generational structure is imposed on the population of evolving individuals, so that every individual is killed off as soon as its fitness is evaluated. In this way, the genetic pool is emptied and then filled up again with newly-reproduced offspring, allowing new combinations of traits to enter the population. In the non-generational evolutionary strategy (Baeck, Hoffmeister, and Schwefel, 1991) and steady-state genetic algorithms (Syswerda, 1989; Whitley, 1989), death does not descend all at once for the individuals in a population, but rather occurs either based on fitness or stochastically. Still, death comes to (almost) all, at a rate determined by the fixed population size and desired rate of introduction of new individuals.² It may be euphemized, as it so often is in everyday language, disguised as "replacement" or "removal," but it is there.

In more ecologically-oriented evolutionary simulations, death – elimination of individuals from the evolving population – may occur for less arbitrary reasons. In such systems, individuals are usually simulated creatures living and behaving in a simple environment. The challenges these creatures face may, if not met correctly, result in the creatures' demise. For example, if a creature lives on energy it amasses as it goes through life, then death can occur if it ever runs out of this energy (as is the case for the system to be described here; see also Ackley and Littman, 1992). Death can also occur at the hands of predators, or by self-inflicted damage (see Ackley and Littman, 1992, again for both of these), which can also be defined in terms of decreasing energy (or health). Combat between individuals is also possible (Holland, 1992, pp. 188-191), as are lethal parasites. All of these forms of mortality, typically not counter-balanced by the introduction of equal numbers of new individuals, imply that the population size will vary, adding to the ecological realism of these simulations.

But what of the creatures that do *not* run out of energy, or die at the hands of predators, or smash themselves against walls or other creatures? In Ackley and Littman's (1992) AL world, no such creatures exist – every individual that begins its life in the world eventually succumbs to its threats. But in other worlds, this will not be the case; as we will describe here, suitably benign simulated environments can lead to the evolution of lucky immortals. Such immortal individuals may even choose to opt out of the reproductive cycle, and thereby retain energy they might have spent in producing offspring, using it instead for their own selfish ongoing survival. When this happens, evolution can grind to a halt, as the immortal creatures fill up the world and prevent the introduction or survival of new individuals. To solve this problem, other forms of death must be introduced to the world, to ensure that evolution can continue.

In this paper, we will first describe an open-ended evolutionary simulation in which the problem of non-reproducing immortals crops up. The appearance of these immortals can follow a curious evolutionary path, as will be seen. This problem, and the associated slowing-down of the evolutionary process, will be addressed with

²In some steady-state GA schemes, the best individuals may float to the top of the fitness-ranked list and be protected there from ever being killed off.

possibilities for new forms of death in the system. These mortality mechanisms will be discussed in terms of their implications for evolutionary simulations, and their relationships to the problems of death and reproduction in real species.

An Open-Ended Evolutionary Simulation

We have begun to develop a framework for exploring the evolution of adaptive behaviors in response to different physical environment structures (Todd & Wilson, 1993) consisting of a simulated world in which a population of behaving creatures evolves. It is in this framework that we will explore the evolution of immortals, and the results of introducing different forms of death into the environment. The environments are simple two-dimensional grids containing food arranged in some layout. The creatures in these worlds can have evolved sensors, internal states, and actions and action-triggering conditions. By allowing all three of these components to evolve, rather than prespecifying any of them, we can explore a wide range of behavior types, including "blind" and memoryless behaviors. Our system is simple and well-defined enough to allow complete specification of the range of possible action-types (including moving, eating, and reproducing) and their effects on the energy levels of the creature and the environment (the bioenergetics of the world).

We now consider the relevant aspects of this system in more detail. We begin with a simple (and easily visualizable) 2-dimensional hexagonal-grid world, each position of which can contain food and/or creatures. Time in the world passes by in discrete clock-ticks, t . The currency in this world is simply energy; creatures live on it and use it up, and food contains it – the greater the amount of food in a particular location, the greater the amount of energy available to a grazing organism.

The world can contain one or more simulated organisms running around in it at a time. Each organism in the world has associated with it a 3-tuple (x,y,E) which codes its current position and internal energy level, and a behavior-generating component which defines its possible actions and triggering conditions, internal state variables, and sensors. The position-energy 3-tuple can change during a creature's lifetime; the beginning values for each creature are determined via inheritance at the time of their birth (or abiogenesis), as will be described later. The behavioral portion of each creature is created by an evolutionary process, and is fixed throughout the creature's lifetime (thus we are for the moment ignoring learning processes).

To explore the evolution of adaptive behaviors, we need a system which can evolve the components of behavior – actions, internal states, and sensors – independently. To this end, we have developed a modular representation of lists of actions and triggering conditions, internal state variables, and sensory input systems. These lists of variables are restricted in how they can interact with each other, so that modularity and interpretability can be maintained. Evolution creates and builds up the entries in each of these lists within each individual creature.

The things that an organism can do in its world depend on what there is in that world to act on and affect. In the system we've described here there is not all that much in the world for the creatures to alter: basically, there is only food, and the creatures themselves. As a result, the number of action-type primitives that we can

define for this world will also be limited. We view this as a plus, though, since it will keep our analysis and understanding of the evolutionary and behavioral dynamics of this system that much simpler.

First of all, creatures can change the distribution of food in the world. The most obvious action in this category is "eating" - lowering the food-value at a particular world location. (This will in turn result in an increase the creature's own energy level, E , and a decrease in the energy of the food, but we will consider these as energetic "side-effects" of the direct action, described in more detail shortly.) For now, we will restrict this class of actions simply to "eat everything that's in your current hex, in one time-step." Another thing in the world that a creature can change is itself. In particular, we currently allow a creature to change its own (x,y) location, moving about in the world by single-hex movement to any of the positions currently adjacent to the creature, making six movement action-types.

There is one final aspect of the world that the creatures should be allowed to influence: the *number* of the creatures themselves. To this end, we add a final action-type, "split." When a creature performs this action, it creates one or more copies of itself as new individuals in the world. (Another action that can change the number of creatures in the world, "die," will be considered later.) We let an individual creature split into two, with each getting an equal amount of energy (to be specified more fully in the world bioenergetics), and each occupying the same hex as the original. (We do not impose any restrictions on creatures occupying the same location in the world grid.) Finally, one of the two new creatures that replace the original is kept identical to that original, and the other is mutated slightly, so that new behavioral mechanisms can enter the population. Without such mutation, the system would be more or less static, with no evolutionary change possible.

As we indicated earlier, every action a creature can perform in its simulated world may have energetic side-effects, in terms of changing the distribution of energy in the world (that is, transferring it from one entity to another). In particular in the current simple world, these effects will be manifested as raising or lowering the individual organism's own personal energy level, E , and changing the energy level in food-plants. How each action-type affects the distribution of energy is determined by the bioenergetics of the world.

The main rule governing the bioenergetics of our system is that *nothing any creature does can raise the energy in the world*. The only way the total world-energy can ever increase is via the food/growth function. Energy can be transferred from one entity to another, and certainly can be lowered or lost, but never increased. If we did not impose such a restriction, then obviously the best thing a creature could do in its lifetime (in terms of survival or ongoing existence) would be to increase its own energy level directly, resulting in essentially a perpetual-motion organism. The only other global rule is that once an individual's energy level E drops to (or below) zero, it dies and is removed from the world and the simulation. This is the traditional form of death in such open-ended evolutionary simulations, and it serves to separate those creatures which are performing more adaptively from those that are performing less so. This form of death thus provides the main force that keeps the simulation in the adaptive-behavior regime we are interested in.

As described above, we have ended up allowing our initial creatures to perform

only three main action-types: eat everything from the current hex, move to one of the adjacent hexes, and split into two. Thus for the energetics of this world, we need only specify the energy side-effects associated with each of these actions. For eating, the energy effects will simply be to lower the individual's energy by an eating-exertion cost (0.5 for now), and then increase the individual's E by the amount of energy absorbed from the food in its current hex and reset the food-energy there to zero (since it is all eaten). The eating-exertion cost guarantees that creatures that always try to eat even when there is no food around will eventually run out of energy and die. For movement, we will charge the creature (i.e., lower its energy E) a movement-exertion cost (1.0 for now) for each step in any direction. For splitting, we impose a splitting-exertion cost (1.0) on the parent creature, and divide the remaining energy equally between the two resulting offspring. Finally, if the creature does nothing at a given time-step, we still impose a small energy sloth-cost on it (0.3 in these runs), so that completely useless sedentary organisms will eventually be weeded out, dying and being removed when their energy slowly leaks away to nothing.

An Example World and the Problem of Immortality

With this framework in place, we began by creating one of the simplest evolutionary scenarios we could imagine: a world filled with food everywhere, populated by creatures without sensors or internal states, only blind actions. In such a situation, all that can be evolved are biases for the performance of certain actions, that is, adaptive rates at which different actions are probabilistically emitted. We wanted to see just what sorts of behavioral repertoires would evolve in such a limited initial case, and how they might be affected by various parameters of the world (e.g. rate of food regrowth, action-exertion costs, mutation rates, etc.).

We constructed a small 10x10 hexagon grid toroidal world, with each edge connected to the opposite one, and filled it with 2 energy-points of food at every location. This food regrows completely at the beginning of the next time-step after it is eaten, so that the world essentially is always full of food. (If there are multiple creatures in a single hex, only one of them will get the food-energy, since it does not regrow *within* a time-step.) The creatures in this world are blind, stateless actors, possessing only an action/trigger list, described in the next paragraph. New creatures are created randomly by abiogenesis (generation of life from inanimate substance) every 100 time-steps, and introduced into the world at the center hex (5,5). Thus competition can build up there, if a pile of immobile creatures remains grazing in that hex; this gives extra impetus for the evolution of movement to escape this tragedy of the commons for greener, less-populated pastures. Each new abiogenerated creature begins life with 10 points of energy.

The creatures' action-lists contain six action/trigger pairs, with the following components: an active/inactive boolean flag, indicating whether or not the following action/trigger is to be used (is active) in the behavioral repertoire of this creature (so creatures can have from 0 to 6 active actions to call on, and inactive actions remain "recessively" available in the creature's genotype); a trigger-value from 1 to 100 which is used to determine how frequently this particular action, if active, will fire (since

there are no sensors or internal states, no more complicated trigger function is possible); and an action type. The possible action types include eat, split, and move.

When we run a world with this set of features, very little happens for some time. New creatures are introduced every 100 time-steps, but they usually die off quickly, running out of energy for one reason or another (basically, because they do not eat often enough in comparison to the other actions they perform). Eventually, a creature may be created that just sits and eats and grows in energy. As long as there is enough food present, such creatures will simply get fatter and fatter and never die.

However, only a certain limited number of these immortal eaters can coexist in the world. In particular, since all the newly abiogenerated individuals are placed in the same hex in the center of the world, as mentioned earlier, competition there for food will be great. With the particular parameters chosen here, 2 energy-units of food per hex per time-step and an eating-exertion cost of 0.5, only three creatures can coexist comfortably in a single hex. If there are four creatures in a hex, then on average they will each only get the 2 units of energy once in every four time-steps, for an average per-time-step energy intake of 0.5, which exactly matches their cost of trying to eat at every time-step. Thus four creatures in a hex will have no net energy change, and will be subject to zero-mean random-walk variations in their energy level which can frequently cause them to run out of energy and hit the absorbing boundary of death. On the other hand, with only three constantly-eating creatures in a hex, each creature will have a net energy gain of 0.166 energy units per time-step ($2.0/3.0 - 0.5$). The creatures' energy levels in this less-crowded case will follow an upward-biased random walk that can buffer them much more effectively from starvation. Three creatures per hex is therefore the carrying-capacity of this world, and if all that ever appeared were immobile immortal eaters, we would only expect to see three of them sharing the center hex.

But as new random creatures are created and dropped into the world, other successful behaviors *do* appear. Some creatures add occasional movement to the mostly-eating repertoire. These individuals will escape the competition in the center hex, as described earlier, and slowly rove around the extent of the world. These creatures can also effectively be immortal: so long as they eat much more often than they move, and there is not too much competition from other creatures, their energy levels can rise without limit. But again the carrying capacity of the world will put a cap on their numbers, still at about three individuals per hex (or 300 for the whole world).

Finally, splitters can appear. We define splitters as those creatures which have the *ability* to split – that is, they have an active split command which could be fired. The most adaptive splitters that arise will eat most of the time, move occasionally, and split to create new offspring very rarely. It is important that splitters move in addition to splitting, so that newly-created offspring will not all end up competing with each other in the same hex. When a splitter like this appears in the world, they can rapidly take over the entire landscape, filling it with their offspring who continue to split and spread further. The resulting population surge gives the splitters domination over their world, and so they reign, but only temporarily....

When we look at a graph of the population characteristics after the arrival of the splitters, a curious finding emerges. As shown in the top line in Figure 1, the population size explodes as soon as the splitters appear (which in this run occurred at a

very early time-step), rises to near the carrying capacity of the world (300 individuals), and stays there, fluctuating slightly. This is not too surprising; one might suppose that the splitters have spread to fill the world, and those that split after the carrying capacity has been reached usually die off, so that the population size is maintained. However, this turns out not to be the case. If we look instead at just the number of splitters in the population, the dashed line in Figure 1, we see that they do rise precipitously in the beginning, driving the population surge, but then their numbers fall off just as dramatically after a few hundred time-steps. By 5000 time-steps, less than 10% of the population still consists of splitters, and their numbers mostly continue to dwindle; by 500,000 time-steps, all of the original splitters have vanished, while the population size still remains near 300. So we see that while the splitters did indeed colonize the world and create the large population which continues to fill it, they do not remain around long to enjoy their conquest. They are quickly eliminated, supplanted by the non-splitters they have given birth to.

Why does this occur? The answer lies in the two remaining curves plotted in Figure 1. These two curves show the average energy of all the creatures in the population (the dashed-dotted line), and the average energy of just the splitters in the population (the dotted line). We put an artificial energy cap of 75 on the creatures in this run, so that any energy gains above this level would be ignored, to keep the energy levels within a restricted range. As can be seen, the population average energy soon rises to near this cutoff value. The average splitter energy, though, remains consistently below this level, meaning that the splitters are that much more susceptible to starvation, as we discussed in terms of random walks earlier. This happens, of course, because every time a splitter splits, it loses half its energy to the other offspring, bringing down the splitter-group's average. Non-splitters never have this problem. In fact, the more often splitters divide, the lower their average energy will fall. This can be seen in the plot of their energy, when compared to the plot of their numbers: when the splitters increase in number (e.g., following time-steps 5000 and 15000), doing so by dividing, their average energy consequently falls. This is clearly an unstable situation, resulting in more dying splitters the more splitters there are. It is no wonder that the splitters fare much worse than their non-splitting brethren, eventually going completely extinct.

The picture that emerges here has faintly sinister overtones. Adaptive splitters appear early on, and eat, move about, and reproduce successfully to fill the world with their progeny. However, some of those offspring have mutated to no longer split, and often not to move either. It is these sedentary eaters which then prosper, outcompeting their still-reproducing parents in the energy game. Eventually, in their single-minded consumption and avoidance of splitting, they alone will survive, and all the hopeful reproducers will starve and die off. The mutants live; the parents die. One could perhaps see this as merely two life-stages in a single species, the first a reproducing, mobile stage, which leads next to an immortal sedentary phase. However, since in the second stage of life these creatures can no longer reproduce, nor change at all themselves, their appearance causes the evolutionary process to grind to a halt, and this is just the problem we must address.

Death, and New Life

The problem facing us in such evolutionary simulations is not really getting rid of the bothersome immortal creatures per se. The behavioral repertoires of such creatures will probably be quite interesting in their own right, and worthy of understanding; but they are not the *only* behavioral responses to the environment that we are interested in. Therefore, what we really want instead is to encourage the evolution of *new* creatures and behaviors. For evolution to be possible, as Kauffman (1991) points out, it is necessary for the system to be able to accrue a succession of small structural changes, each with an associated small change in fitness. Abiogenesis - peppering wholly-new random creatures into the world - obviously cannot do this (it constitutes a random search of behavior-space, rather than the sort of gradual quasi-hill-climbing that evolution can perform). Only reproduction, here splitting to create a new slightly-altered individual, can. Thus, to encourage evolution in our system, we must encourage reproduction. And to encourage reproduction, what better than looming death?

Waddington (1982) states that "[d]eath is a necessary condition for the trying out of new genetic combinations in later generations" (p. 649). There are two problems with this conjecture as regards natural populations. First, as the Waddington quotation at the beginning of this paper made clear, it is true only if the environmental niche that all offspring are born into is completely full, so that they cannot compete with the already-established individuals and so soon die. If the niche is *not* at capacity, then newly-born creatures can hope to live without older individuals dying off to make room; and if the new creatures are different enough from their parents (through mutations) that they seek to colonize a *new* econiche, they will escape the competition of their elders. Second, the implication here is that evolution "wants" to try out "new genetic combinations in later generations", but of course this is not the case. The process of evolution may cease with the absence of such trials, but since evolution has no goal - including its own continuation - it will not create death as a way to keep going. There is no logical reason why evolution could not create an immortal species, and then stop (for that species at least).

But when applied to evolutionary simulations, rather than the real world, Waddington's statement rings more nearly true. In our simulation, with a limited amount of energy to be shared among all the creatures in a very limited world, it is easy to arrive at an overcrowded situation in which no newly-born creature could survive. This is indeed what we saw in the example in the previous section, when the creatures reached the carrying capacity of the world. Thus, death - the removal of existing creatures - probably *is* required for new individuals to be "tried out". Furthermore, generally in our simulations we *do* have a goal for evolution - we want it to come up with the most interesting and lifelike creatures possible, and to continue doing so as long as possible. (One quantitative way to measure the lifelike activity in the system could be to look at the amount of energy-transfer between entities taking place: worlds with few creatures, or those that just sit and eat, will have little energy flow, while those with lots of active creatures, incorporating and expending energy and reproducing, will show greater energy flow.) We don't want it to just stop at some point, when there are always further potential improvements. So we do demand "the

trying out of new genetic combinations", and thus we demand death.

But how shall this sentence be imposed? There are a few possibilities. We already have death due to starvation (running out of energy), but as the existence of immortals in the simulation showed, this is not sufficient. We must also introduce a mortal force from which escape is impossible. The most obvious is a fixed lifetime, killing off all creatures that reach a certain ripe old age. This could also be implemented stochastically, with an increasing rate of mortality with increasing age. Still more stochastically, we could have all individuals subject to some rate of demise no matter what their age. This would be like hailing down a storm of meteorites with a certain frequency onto the world, killing off creatures randomly with some (low) probability at every time-step. All three of these methods would be essentially inescapable - nothing a creature could do would keep it alive indefinitely, and all are likely to die within some determinable period of time.

In the face of this impending doom, the most adaptive course of action is to reproduce, to extend one's genetic representation into the future. We hypothesized that, the shorter the effective induced lifespan (whether by fixing it at some number of timesteps, or imposing it stochastically), the quicker our evolved creatures would reproduce - that is, the more frequently the split action would be chosen. In initial simulations, this has indeed been the case (we are still collecting data on this score). This of course only applies to those creatures that reproduce at all; after such splitters have evolved for a while, tuning their behavioral repertoires in terms of firing rates, the inverse relationship between lifespan and frequency of reproducing appears.

Another less satisfactory means of fatality is to kill off all those creatures that reach a certain energy-level (rather than just capping their energy there, as we did in this simulation). Such bloated creatures could be said to explode because of eating too much, for instance. But such a form of death is not inescapable, and we hypothesize that creatures will evolve to keep their energy-level near but below this lethal limit, through a carefully-crafted combination of energy-gaining (eating) and -expending (moving and/or splitting) actions. (Note that if we had imposed a lethal energy limit of 75 points in the simulation described earlier, the splitters would have mostly succeeded in avoiding it; but non-splitting immortal creatures could also have escaped, which is why this killing method is not sufficient.)

These and other externally-imposed causes of death may serve to solve the problem of non-reproducing immortals, necessitating reproduction for adaptive species and ensuring the continuation of the evolutionary process in our simulations. They assume an inevitable senescence on the part of our creatures, such as an accrual of metabolic toxins or a breakdown of repair mechanisms that occurs with advancing age and eventually results in total system failure (Comfort, 1979). But such physiological processes are not the only possible causes of senescence. Evolution may also have programmed living species for senescence and death at certain ages through the course of natural selection. Indeed, Bell (1984) has collected evidence that this is the more likely cause of senescence. Such evolved senescence, or post-reproductive decline in vitality, could arise because of the accumulation of deleterious genes that negatively effect health in later age, and which natural selection is less effective at weeding out (Medawar, 1952). Alternatively, genes selected for their positive effect on fitness early in life, could have a negative effect on viability later in life (Williams, 1957). Rose

and Charlesworth (1980) have shown experimental support for the latter theory, and Partridge and Barton (1993) review the evidence for both. In still another much more limited domain of death, de Catanzaro (1987) has proposed an evolutionary theory of suicide, based on an individual's self-assessment of its own fitness prospects versus those of its kin.

These evolutionary issues are all topics we hope to address with our simulation models, since they are at the heart of a discussion of the hows and whys of death. To tackle these questions, though, we must modify our models in some significant ways. We cannot begin to look at the evolution of senescence, for instance, until we include a genetically-controlled developmental process, or at least life-history, in our models. Only then will we be able to model the effects of genes that cause changes in viability at various stages of life, as the first two theories just mentioned require. In investigating de Catanzaro's theory, we can now simply introduce a "die" or "suicide" action, which if performed causes the creature to fall immediately dead. If the creature's energy is then simply removed from the world, such an action will probably not prove useful to anyone, and indeed we see this form of suicide weeded out and suppressed in our simulations. However, if the suicidal creature's energy is transferred to the food in its final resting-hex, available to be eaten by other creatures, then we may see this type of behavior spread in the population in cases where this energy "inheritance" benefits the offspring or relatives of the deceased. We are currently exploring this situation to see whether even our simple sensor-less creatures can take advantage of it, or if more complicated mechanisms of kin assessment will be necessary before this type of suicide can possibly be adaptive. Thus, many questions regarding the evolution of death remain to be investigated with further enhancements of our model.

Conclusions

But for now, our explorations have begun to show the usefulness of artificial death in evolutionary simulations. Without imposed fatalities, non-reproducing immortals can evolve that choke off reproducing lineages and result in an end to the evolutionary process. With inescapable forms of mortality, though, successful (adaptive) creatures are forced to resort to reproduction to continue their genetic representation into the future. The more imminent death is, the sooner these creatures will choose to reproduce. In this way, by increasing the death-rate, the evolutionary "drive" to create new offspring can be bolstered, thereby potentially speeding up our simulations. In continuing research, we will further explore the relationship between death and reproduction, and begin to investigate some of the theories of the evolution of death and senescence described in the previous section. In the meantime, we hope to have shown that, at least in the realm of evolutionary simulations, out of the jaws of death some benefit may come.

References

- Ackley, D., and Littman, M. (1992). Interactions between evolution and learning. In C.G. Langton, C. Taylor, J.D. Farmer, and S. Rasmussen (Eds.), *Artificial life II* (pp. 487-509). Redwood City, CA: Addison-Wesley.
- Baeck, T., Hoffmeister, F., and Schwefel, H.-P. (1991). A survey of evolution strategies. In R.K. Belew and L.B. Booker (Eds.), *Proceedings of the Fourth International Conference on Genetic Algorithms* (pp. 2-9). San Mateo, CA: Morgan Kaufmann.
- Bell, G. (1984). Evolutionary and nonevolutionary theories of senescence. *American Naturalist*, 124, 600-603.
- Comfort, A. (1979). *The biology of senescence* (3rd edition). New York: Elsevier.
- de Catanzaro, D. (1987). Evolutionary pressures and limitations to self-preservation. In C. Crawford, M. Smith, and D Krebs (Eds.), *Sociobiology and psychology: Ideas, issues, and applications* (p. 311-333). Hillsdale, N.J.: Lawrence Erlbaum Associates.
- Goldberg, D.E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Reading, MA: Addison-Wesley.
- Holland, J.H. (1992). *Adaptation in natural and artificial systems*. Cambridge, MA: MIT Press/Bradford Books. (Expanded edition of 1975 original published by University of Michigan Press.)
- Kauffman, S.A. (1991). Requirements for evolvability in complex systems: Orderly dynamics and frozen components. In S. Forrest (Ed.), *Emergent computation: Self-organizing, collective, and cooperative phenomena in natural and artificial computing networks* (pp. 135-152). Cambridge, MA: MIT Press.
- Malthus, T.R. (1992). *An essay on the principle of population, or, A view of its past and present effects on human happiness : with an inquiry into our prospects respecting the future removal or mitigation of the evils which it occasions*. Cambridge: Cambridge University Press. (A new edition of the 1803 original.)
- Medawar, P.B. (1952). *An unsolved problem of biology*. London: H.K. Lewis.
- Partridge, L., and Barton, N.H. (1993). Optimality, mutation and the evolution of ageing. *Nature*, 362, 305-311.
- Rose, M., and Charlesworth, B. (1980). A test of evolutionary theories of senescence. *Nature*, 287, 141-142.
- Syswerda, G. (1989). Uniform crossover in genetic algorithms. In J.D. Schaffer (Ed.), *Proceedings of the Third International Conference on Genetic Algorithms* (pp. 2-9). San Mateo, CA: Morgan Kaufmann.
- Todd, P.M., and Wilson, S.W. (1993). Environment structure and adaptive behavior from the ground up. In J.A. Meyer, H.L. Roitblat, and S.W. Wilson (Eds.), *From animals to animats 2: Proceedings of the Second International*

Conference on Simulation of Adaptive Behavior (pp. 11-20). Cambridge, MA: MIT Press/Bradford Books.

Waddington, C.H. (1982). Development, biological. In *The new Encyclopaedia Britannica*, vol. 5 (pp. 643-650). Chicago: Encyclopaedia Britannica Inc.

Whitley, D. (1989). The GENITOR algorithm and selection pressure: Why rank-based allocation of reproductive trials is best. In J.D. Schaffer (Ed.), *Proceedings of the Third International Conference on Genetic Algorithms* (pp. 116-121). San Mateo, CA: Morgan Kaufmann.

Williams, G.C. (1957). Pleiotropy, natural selection and the evolution of senescence. *Evolution*, 11, 398-411.




Figure 1. The population size (solid line) and average energy of all creatures (dashed-dotted line), and the subpopulation size of splitter creatures (dashed line) and their average energy (dotted line), showing the surge and decline of splitters.