

Young Scientists' Corner

Natural Selection as an Algorithm: Why Darwinian Processes Lack the Information Necessary to Evolve Complex Life

Natural Selection as an Algorithm: Why Darwinian Processes Lack the Information Necessary to Evolve Complex Life



The invention of computers and the advent of computer science ... have begun to cast doubt upon Darwin's vision of a bottom-up world in which simplicity gives rise to complexity.

Darwin's great insight was his explanation of the complex in terms of the simple. His elegant mechanism, natural selection, relies upon the facts of replication, variation, and competition to explain high-level complexity and even the very intelligence capable of understanding and contemplating evolution itself. The capabilities of the Darwinian mechanism appear to be literally limitless, as it creates and designs the well-wrought, complex life that inhabits the world.

The invention of computers and the advent of computer science, however, have begun to cast doubt upon Darwin's vision of a bottom-up world in which simplicity gives rise to complexity. The computer is an ideal modeling tool for the slow, drawn-out process of biological evolution, since many millions of operations can be carried out per second. One of the first to apply the budding field of computer science to evolutionary theory in the 1960s was John Holland of the University of Michigan, who coined the phrase "genetic algorithm" and, along with his students, began to experiment with creating a computer model of the evolutionary process. As the field of evolutionary programming has grown, it has begun to demonstrate that the complexity produced by genetic algorithms first must be encoded into the parameters and fitness functions—it does not arise from scratch.

To illustrate this, consider the public lecture that first sparked my interest in genetic

John Bracht received a B.S. in biology from New Mexico Tech and is beginning doctoral studies in cell biology. Following his interest in complex systems and design theory, John is a founding member of the International Society for Complexity, Information, and Design (<http://www.iscid.org>), and has published in the society's journal, *Progress in Complexity, Information, and Design*. John has also written for the online science and religion forum *Metanexus*, for the *Templeton Foundation's* *Research News and Opportunities in Science and Theology*, and the *Santa Fe Institute's* *Journal, Complexity* (in press). Send correspondence to him at: jbracht@ucsd.edu

algorithms. The seminar was presented on February 21, 2001, at New Mexico Tech in Socorro, NM, while I was an undergraduate student there. Dave Thomas, an alumnus of New Mexico Tech and current president of the local skeptics group, New Mexicans for Science and Reason, presented a computer program that purportedly generated specified complexity entirely without intelligent input or design. The program's task was to solve the Steiner problem, which entails finding the minimal network that connects five pre-given points. The program began by generating a series of random networks. Then through a series of random "mutations" and rounds of selection, the program was able to converge to the optimum Steiner solution with great regularity.

After his presentation, Thomas and I had a lengthy email exchange in which I pointed out that the Steiner solution is the network that connects all five points and has the shortest path-length. But the genetic algorithm selected for networks that connect all five points and have shortest path-lengths. Thus, the very properties that define the Steiner solution were programmed into the fitness function. Notice how specific the fitness function had to be: it captured the defining characteristics of the desired solution. The high degree of specificity embodied in the fitness function also prevented the program from ever converging on any other configuration that might be useful or interesting, or from providing a solution to a different problem; in that sense, the program was deterministic. The implication is that there are no general-purpose genetic algorithms. They must be carefully tuned to fit the problems they are supposed to solve, and they must contain large amounts of very detailed information.

Several prominent scientists in the field of genetic algorithm theory echo this conclu-

sion. Geoffrey Miller of University College, London, notes that fitness functions must be carefully designed. He says:

In effect, the fitness function must embody not only the engineer's conscious goals, but also her common sense. This common sense is largely intuitive and unconscious, so is hard to formalize into an explicit fitness function. Since genetic algorithm solutions are only as good as the fitness functions used to evolve them, careful development of appropriate fitness functions embodying all relevant design constraints, trade-offs and criteria is a key step in evolutionary engineering.¹

Furthermore, Miller notes:

All the expertise that human engineers would use in confronting a design problem—their knowledge base, engineering principles, analysis tools, invention heuristics and common sense—must be built into the genetic algorithm. Just as there is no general-purpose engineer, there is no general-purpose genetic algorithm.²

Melanie Mitchell, expert on genetic algorithms from the Santa Fe Institute and author of an introductory text on genetic algorithms, notes:

Choosing a fixed encoding ahead of time presents a paradox to the potential GA [Genetic Algorithm] user: for any problem that is hard enough that one would want to use a GA, one doesn't know enough about the problem ahead of time to come up with the best encoding for the GA. In fact, coming up with the best encoding is almost tantamount to solving the problem itself!³

In other words, the amount of design work and information a genetic algorithm needs to solve a problem is often enough to permit a direct solution, without the genetic algorithm as an intermediate step. The important point is that the particulars of the problem and the desired outcome both must be explicitly put into the system before it can solve the problem. The genetic algorithm, especially the fitness function, is highly specific and contains very detailed information about what it is to produce. The complexity and information is all there in the antecedent conditions before the program begins to operate, and the program acts as a conduit for that information to flow into the end result. The information content of the resulting simulated organisms does not originate step-by-step; rather, it is present, in its entirety, in the program at the beginning, and it is incorporated into the end result in a step-by-step fashion.

Another example will help illustrate this point. Thomas Schneider, research biologist in the Laboratory of Experimental and Computational Biology at the National Cancer Institute, has published an article in *Nucleic Acids Research* in which he describes a genetic algorithm he created.⁴ His

program generates a population of 64 "organisms" each having a genome of 256 "base pairs." These base pairs may take one of four values (to simulate the four bases available to real DNA). The goal of the simulation is to model the evolution of binding sites in the genome of hypothetical organisms. Schneider set up sixteen 6-base long binding sites and a weight matrix that is used to assign numerical values to each binding site. In addition, the genome contains an encoded threshold value, and if the numerical value of a binding site (as determined by the weight matrix) is above the threshold, a hypothetical protein is considered to have bound to the site.

The complexity and information is all there in the antecedent conditions before the program begins to operate, and the program acts as a conduit for that information to flow into the end result.

The fitness function simply counts the number of mistakes made by each organism and eliminates the half of the population that make the most mistakes. Those making the fewest mistakes then replicate (with one mutation per organism) and replace those organisms that have been eliminated. A mistake is defined as non-binding at a binding site, or binding at a non-binding site. In other words, if the numerical value of a binding site is below threshold, or the numerical value of a non-binding site is above threshold, a mistake has been made. What this implies is that the fitness function selects directly for organisms with binding site values above threshold (and no non-binding site areas above threshold).

Schneider comments: "Remarkably, the cyclic mutation and selection process leads to an organism that makes no mistakes in only 704 generations." Furthermore, he notes:

The ev model quantitatively addresses the question of how life gains information, a valid issue recently raised by creationists ... The ev model shows explicitly how this information gain comes about from mutation and selection, without any other external influence, thereby completely answering the creationists.⁵

However, it is abundantly clear from the analysis given above that the information produced by the program was actually smuggled in by the programmer. The fitness function selects directly for organisms that have the most binding site values above threshold and have the fewest non-binding site values above threshold. Is it any wonder, then, that the program produces organisms that have all

Young Scientists' Corner

Natural Selection as an Algorithm: Why Darwinian Processes Lack the Information Necessary to Evolve Complex Life

binding site values above threshold and have no non-binding site values above threshold? Again what we see is that the very characteristics of the desired outcome are put directly into the fitness function, and there is no magic or mystery about where this information comes from. It is not generated "from mutation and selection, without any other external influence." It is inputted by Schneider himself and is extracted and made explicit by the mutation and selection process.

The program Schneider produced has several characteristics which cannot be generated by selection and mutation and which are necessary before selection and mutation may begin to operate. These characteristics manifest as certain fixed parameters that Schneider prepares ahead of time. They constitute the *encoding* of the problem—the settings that configure the program to deal with things like simulated DNA sequences and binding sites.

The fact that mutation has to have something to work upon requires some sort of genome and system of translation or expression, which are given by the weight matrix and binding sites. Even though the binding sites may change position between runs, the precise position of binding sites is fixed within a run, and the computer program must know where the binding sites are in order to know how to evaluate the number of mistakes (recall that whether a site is a binding-site is essential to whether binding is desirable). Furthermore, the fitness function must have a logical structure that allows it to determine which sites to favor and which to reject. This entails the complex interrelationship among weight matrix, threshold, and interpretation of binding site value that Schneider set up ahead of time. These parameters, the encoding, are absolutely fixed but could conceivably take any number of possible configurations—the options are as limitless as the number of possible problems that could be encoded into a computer program.

The selection of one option from many in the encoding amounts to the introduction of huge quantities of information, and essentially gets the program "in the ballpark" for finding the answer Schneider wants. The fitness function is overlaid upon the encoding

parameters, and functions to direct the program to the right solution within the pre-given encoding. The fitness function, then, is something like a "map" of the ballpark, showing the program precisely where to go to find the solution (or, more precisely, defining what a solution *is*). Therefore, to claim (as Schneider does) that his program generates information "from scratch" and requires no intervening intelligence is patently false. The information was carefully encoded by the choice of inserted and fitness function to allow something interesting, like the evolution of binding sites, to happen.

These examples show that genetic algorithms always must start with the information they output. The fitness function must select directly for the defining qualities of the desired outcome. In other words, evolutionary processes precisely follow their fitness function. If the fitness function does not go anywhere, neither does the evolutionary process. If the fitness function explicitly points toward some complex and interesting result, then that is what the evolutionary process will produce. The question we need to consider, then, is: Does biological selection point toward complex and interesting organisms like ourselves? Does it contain the sort of explicit, detailed information needed to drive an evolutionary process to produce complex, higher organisms? Even a cursory examination of natural selection is enough to show that it does not point toward anything complex and interesting, and indeed that it should penalize increases in complexity.

It is commonly observed that evolution has no goal, no purpose in mind. It is a process which relies on nothing more than the sifting of variants; a form of filtered randomness. Schneider's program starts with an end goal, a perfect picture of what it is to work toward. In fact, this is a general characteristic of genetic algorithms: they require a detailed picture of what they are to produce before they can produce it. They cannot simply generate random things and grab "whatever looks interesting." They require a specific, detailed idea of "what is interesting," a goal that they can work toward. The problem is that the Darwinian fitness function, natural selection, *has no analogous goal*. It does not have an internal "human-producing" module that contains complete specifications for a human being. Yet all of our experience in

Genetic algorithms always must start with the information they output. The fitness function must select directly for the defining qualities of the desired outcome. In other words, evolutionary processes precisely follow their fitness function.

genetic algorithms suggests that the only way evolution could produce complex entities like human beings is if that endpoint is specifically targeted by the fitness function— that is, if humans (and all the rest of biological complexity) are explicitly identified as targets before evolution begins to work. If this is true, then the sheer volume of information that must be contained within the Darwinian fitness function must be enormous, and we are justified in asking where all of this information is stored.

By invoking the Darwinian mechanism as the cause of all biological complexity, [Darwinists] imply that somehow, without any goal-directedness, their mechanism was able to sift randomness and preserve “whatever is interesting” – all without knowing, ahead of time, what the definition of “interesting” is.

Darwinists are quite right to insist that their mechanism actually does not contain pre-encoded endpoints, or goals; it does not contain the vast amount of information needed to precisely describe all of biological complexity. It is here, though, that a problem arises. For by invoking the Darwinian mechanism as the cause of all biological complexity, they imply that somehow, without any goal-directedness, their mechanism was able to sift randomness and preserve “whatever is interesting” – all without knowing, ahead of time, what the definition of “interesting” is. And as we have seen, computer simulations of the evolutionary process demonstrate that only a precisely defined fitness function, with a detailed goal in mind, can produce complex and interesting things. In short, the program has to know exactly what end goal to select for – and evolution simply does not have any knowledge of such goals. Thus, in a detailed, goal-oriented sense, evolution is *in principle* incapable of producing biological complexity.

This is where the plot thickens. The Darwinian argument is that evolution need not have an end goal in mind because it can substitute “fittest” for “interesting.” Evolution can just generate random variations of each organism, and, by preserving (selecting) the fittest variant (defined as those organisms capable of most efficiently producing offspring) is able to boot-strap itself up to build all biological complexity (given enough iterations of the cycle). In that sense, the end goal of the Darwinian process is not any particular piece of biological complexity, but rather a good-enough substitute: the fittest organism.

There are now two problems facing the Darwinian mechanism, and I will focus here upon the second (the first has been dealt with elsewhere by myself and other authors).⁶ The first problem is that generating new variants to test (new candidates to run through the filter of selection) is the job solely of random mutation, and those mutations must produce new adaptations which natural selection can then preserve. This works fine for certain types of adaptations (like antibiotic resistance or finch beak size evolution) but there are some adaptations which require too big a step for chance to take. The reason is that a novel adaptation sometimes requires multiple, coordinated changes in many genes (or the origin of many entirely new genes) before any selectable advantage is realized, putting such adaptations simply beyond the reach of a chance-based adaptation generator. Because of the interlocking requirements for function, there is no gradual route to such adaptations; they must be formed in a single step – a step that is beyond the reach of chance. And if chance cannot generate those adaptations, they will never exist for natural selection to act upon and these adaptations will never be generated via a Darwinian mechanism – yet many such adaptations actually do exist in the biological world.

The second problem concerns the Darwinist’s proposed substitution of “fittest” for “interesting” as a goal for the evolutionary process. By “interesting” I mean biological complexity beyond simple bacteria and other microorganisms (we could call this “higher biological complexity,” and it includes ourselves). “Fitness,” in a biological sense, is a measure of one’s ability to propagate oneself. I intend to challenge the link between “fittest” and “interesting.” There is no reason they should be synonymous, and there are good reasons to suspect that in most or perhaps all cases they are actually antonymous.

Let’s take another look at natural selection. Richard Dawkins, in *The Blind Watchmaker*, describes how natural selection operates to increase the fitness of organisms in a population. He says:

And if any entity, anywhere in the universe, happens to have the property of being good at making more copies of itself, then automatically more and more copies of that entity *will* obviously come into existence. Not only that but, since they automatically form lineages and are occasionally miscopied, later versions tend to be “better” at making copies of themselves than earlier versions, because of the powerful processes of cumulative selection. It is all utterly simple and automatic. It is so predictable as to be almost inevitable.⁷

The problem here is that “fit” (efficiently replicating organisms) and “interesting” (higher biological complexity) are diametrically opposed to each other. As the complexity of a system increases, so does the cost (in time

Young Scientists' Corner

Natural Selection as an Algorithm: Why Darwinian Processes Lack the Information Necessary to Evolve Complex Life

Natural selection favors simple, single-celled replicators and penalizes higher complexity and organization as we see in multicellular life and in higher animals. The Darwinian world should consist entirely of bacteria and other single-cellular life forms competing to be the simplest and most efficient replicator.

and material resources) of making more copies of it. Therefore, the best way to increase reproduction efficiency is to reduce complexity. In other words, the selection pressures of evolution point *away* from the advanced biological complexity that we consider “interesting.” Richard Dawkins addresses this point when he describes Sol Spiegelman’s experiments in the 1960s with RNA replication in test tubes.⁸ The RNA used in the experiment originally came from a Q-beta virus (which normally infects *E. coli*), and it encodes an RNA replicase used by the virus to duplicate its RNA and take over a cell.

Spiegelman prepared a test tube with pre-made RNA replicase, and added the Q-beta RNA. After allowing the RNA to go through replication, he took a drop from that test tube and added it to a fresh test tube (again, with pre-made replicase present). The result was an unequivocal drive toward simplicity, with the most successful RNA replicators being those molecules which managed to jettison the unneeded RNA replicase sequence. Since Spiegelman was supplying replicase in high quantities, there was no need for the viral RNA to carry instructions for making more—and the RNA was able to replicate much faster by becoming as simple as possible, in this case, simplifying down to the bare minimum needed to carry out the act of copying. Even some original complexity in this relatively simple replicator was lost since it was not absolutely essential to the replication process.

The clear implication is that natural selection favors simple, single-celled replicators and penalizes higher complexity and organization as we see in multicellular life and in higher animals. The Darwinian world should consist entirely of bacteria and other single-cellular life forms competing to be the simplest and most efficient replicator. After all, bacteria, which can reproduce (under optimum conditions) every twenty minutes, far outstrip the twenty-year life cycle of humans in terms of replication efficiency. In general, as one moves up the complexity ladder, one finds a corresponding increase in life cycle time and a consequent decrease in replication efficiency. Furthermore, bacteria are far more adaptable than humans and are found in every environment that humans inhabit—

and quite a few that we do not (just think of the deep thermal ocean vents).

The increased complexity of human beings (or any other higher life form) gives no benefit in terms of habitable environment or fitness in a Darwinian sense. In fact, it is often noted that if some global catastrophe occurs (such as a nuclear exchange), the only organisms that survive will be the bacteria and relatively simple organisms such as cockroaches. This only highlights the fitness cost (in terms of survival, not just reproduction) of increased complexity. It is difficult to see what possible benefit could accrue from increased complexity such as to overcome the corresponding decrease in replication efficiency and overall fitness. The fittest, in a Darwinian sense, are the bacteria and other micro-organisms of which we are hardly aware. Humans and other multi-cellular organisms are anomalies—lumbering, gigantic, and ponderously unfit in the Darwinian world.

The brute fact of the existence of beings like ourselves suggests one of two things; either (1) we were not “in the program” of Darwinian evolution (and hence came from a non-Darwinian process) or (2) we are “in the program” of evolution (and need to find out where). Let us consider the implications of the second possibility. Perhaps the Darwinian fitness function is, in reality, more complicated and nuanced than I am giving it credit. Perhaps scientists will one day find that it does contain the sort of information required for complex life like ourselves to arise; perhaps the evolutionary fitness function points directly and unequivocally toward complex organisms. Precisely what would this evolutionary fitness function look like?

Before we consider the information required to make a human being, let us consider the much simpler case of a “mere” bacterial flagellum. In this case, we have a good idea of the constraints and requirements that would be needed to create this miniature outboard motor, since extensive research has been done on the system.⁹ We know that it normally rotates at 20,000 RPM, that it has various rotors and stators, a “propeller” and hook joint (a sort of universal joint on a molecular scale), and an elegant system for converting acid flow to rotary

motion. Recall that the study of genetic algorithms reveals that fitness functions must select for *the defining characteristics of the desired result*. Therefore, if natural selection were to produce a complex system like the bacterial flagellum, it would have to somehow select for a rotary motor, complete with drive shaft, propeller, and acid-powered engine. Granted, there may be more than one way to construct such a motor, but such possibilities are very limited considering the vast space of possible configurations of protein molecules. It is to this highly specialized target area that natural selection must guide the evolutionary process if it is to create a bacterial flagellum. Under these conditions, a flagellum should be expected to emerge, and always emerge. However, this fitness function would be unable to select for anything *but* a bacterial flagellum. We run into the determinacy problem of genetic algorithms; they always converge on their target, with a probability of one.

The complex design produced by the algorithm must be programmed into the fitness function from the outset. In defiance of Darwin's vision of a bottom-up, step-by-step route to complex life forms, genetic algorithms are demonstrating that the complexity and order inherent in life is not reducible to simpler components.

This flagellum-building fitness function certainly could not work to build other complex structures like the blood clotting cascade or eukaryotic cilium, let alone the brain/eye system or the intricacies of the circulatory/respiratory system required by large organisms like ourselves. The problem here is that the evolutionary algorithm is too specific and cannot function as a universal problem solver to produce all the different types of order in the biological world. Even given enough information to produce some sort of complex life, it would lack the ability to produce other sorts of complex life.

Perhaps even more importantly, there is absolutely no reason to suspect that the evolutionary fitness function does anything like selecting for bacterial flagellum proteins; in nature, we only observe selection for reproduction. Indeed, if the evolutionary process did select for proteins useful to making a bacterial flagellum it would be deeply teleological, working toward an overall goal—and the oft-

cited benefit of Darwin's idea is that it separates teleology from (apparent) design. Thus, even if we eventually find that a teleological fitness function is responsible for some or all complex life in biology (an idea unsupported by current knowledge), it would be a profoundly non-Darwinian mechanism.

The bottom line is that an evolutionary fitness function sufficient to do the design work of complex life forms would itself have to be designed and very complex. Also, it would have to be re-designed for each new feature that emerged—there is no universal genetic algorithm which can create a bacterial flagellum one moment and then build a vertebrate eye the next. These multiple fitness functions would each embody the design and complexity they create, and thus they would simply beg explanation, making the design problem in biology even more acute. However, such elaborate fitness functions do not exist in nature, at least as far as we can tell, and the central claim of Darwinism is that we do not need them—that we can explain the complexity of life in terms of the simple fitness function of natural selection. But computer simulations of evolution have shown the inadequacies of such a simplistic model. There is no universal problem-solver, and each fitness function must be carefully tuned to select for the desired outcome. The complex design produced by the algorithm must be programmed into the fitness function from the outset. In defiance of Darwin's vision of a bottom-up, step-by-step route to complex life forms, genetic algorithms are demonstrating that the complexity and order inherent in life is not reducible to simpler components. The complexity and design of life is holistic; it is top-down, not bottom-up. And that concept is profoundly non-Darwinian. ★

Notes

¹Geoffrey Miller, "Technological Evolution as Self-Fulfilling Prophecy," in *Technological Evolution as an Evolutionary Process*, ed. John Ziman (Cambridge: Cambridge University Press, 2000), 207.

²*Ibid.*, 209.

³Melanie Mitchell, *An Introduction to Genetic Algorithms* (Cambridge: MIT press, 1998), 158.

⁴Thomas Schneider, *Nucleic Acids Research* 28, no. 14 (2000): 2794–9.

⁵*Ibid.*

⁶See chap. 5 of *No Free Lunch* by William Dembski (Lanham: Rowan & Littlefield, 2002). My own writings on the problems for a Darwinian mechanism posed by the bacterial flagellum can be found on Metanexus (<http://www.metanexus.net/index.html>). See my response to Matt Young (found at www.metanexus.net/archives/message_fs.asp?list=views&listtype=Magazine&action=sp_simple_archive_&page=1&ARCHIVEID=6385&searchstring=Bracht) and forthcoming response to Ursula Goodenough dealing extensively with the bacterial flagellum.

⁷Richard Dawkins, *The Blind Watchmaker* (New York: Norton, 1987), 134.

⁸*Ibid.*, 300–4.

⁹See, for instance, Michael Behe, *Darwin's Black Box* (New York: Simon & Schuster, 1996), 70–2.