What Does Evolutionary Computing Say About Intelligent Design?

Karl D. Stephan

1. Introduction

When Charles Darwin proposed his theory of variation and natural selection in *The Origin of Species*, he unwittingly provided the intellectual foundation for what has recently become a thriving subfield of computer science. Evolutionary computing arose from artificial intelligence research in the 1960s by computer scientists such as Lawrence Fogel and John Holland. Holland was one of the first to show that the genetic mechanisms of variation, reproduction, and selection are peculiarly suited to deal with certain problems whose characteristics make them difficult to solve with conventional approaches (Holland, 1975; Holland, 1992). In the decades since, dozens of books and hundreds of papers have been published on the subject, ranging from theoretical advances in the genetic algorithms used in evolutionary computing to practical applications in molecular modeling, mechanical and electrical engineering, and even investment strategies.[1] Evolutionary computing today is much more than a twinkle in the eyes of a few computer scientists. In its wide range of successful applications, it is one of the relatively few concepts in artificial intelligence which has actually lived up to the hopes of its early proponents.

Does the practical success of evolutionary computing say anything about the theory of natural evolution? Beyond that, we can ask a more specific question: does evolutionary computing give support to the contention that variation and natural selection alone are inadequate to explain the origin and variety of life? This contention is often referred to as the “intelligent design” hypothesis, since it denies that nature, unaided by intelligence, is capable of producing the plethora of present life forms on Earth. Before we begin to answer either of these questions, we shall first describe evolutionary computing in enough detail to elucidate some of its basic approaches. Then we shall point out the parallels and differences between the trajectory of information in evolutionary computation, and the way information is stored, modified, and used in in natural biological systems such as bacteria. At that point, we will be prepared to address the two questions above.

2. What Is Evolutionary Computing?

Evolutionary computing was developed by computer scientists to solve difficult problems. Historically, one of the main driving forces behind the development of electronic computers was the desire on the part of scientists and engineers to solve problems that were too complex to solve any other way. For example, shortly after World War II, the complicated physics of nuclear weapons at high pressures and temperatures was simply too much to handle with conventional methods that used pencil, paper, and mechanical calculating machines. The dangers and expense of building an untried bomb design just to see if it works are obvious. So numerical simulation of nuclear explosions became one of the first important uses of electronic computers.

This example brings up a point of critical importance to this discussion: the distinction between analysis and synthesis in design. Analysis answers the question, "If we make such-and-such a structure under such-and-such conditions, what will happen?" Analysis "takes apart" a situation and breaks it down into simple pieces so that mathematical and physical laws can be applied to predict the system's behavior. Nature (and in particular, quantum mechanics) being what it is, no analysis can make an absolutely perfect prediction about a real-world system down to every last detail. But as the speed and capacity of computers have increased, the scale and accuracy of computer-based analysis have grown to encompass larger and more complex systems, from the steel structures of entire skyscrapers to models of global climate.

At first, designers used computers only for analysis. Synthesis–the job of originating a design to be analyzed–was still a task of the human mind. But it was not long before scientists began to instruct computers to vary designs which were initially provided to them by a human designer. The computer then analyzed the effect of each variation on the performance of the system under study, and stopped when it found variations that produced an "optimum" design, in some sense.

The computational procedures for doing this sort of thing became known as optimization algorithms. They did not always work, most often because the problem under analysis either did not have a well-defined optimum, or was otherwise ill-behaved. But in certain specialized areas in engineering and physics, computer optimization proved to be very useful.

In computer optimization, the computer does a little synthesis as well as analysis. In trying its variations to seek an optimum solution, the computer "originates" designs that the human designer did not explicitly call for. But typically, the human designer sets the limits over which the computer can try its own variations. In this sense, the computer is on a very short leash. It can try certain novel things within well-defined ranges, but cannot stray outside those ranges.

Evolutionary computing uses procedures that overcome many of the disadvantages that conventional optimization algorithms have. As the name implies, evolutionary computing uses concepts borrowed from Darwinian evolution to "breed" good solutions to design problems. Although the details of how evolutionary computations work can be complicated, the basic ideas are simple. One standard example is the way evolutionary computation expert David Fogel (son of Lawrence Fogel) solves the "traveling salesman" problem (Fogel, 2000).

The traveling salesman problem is easy to state. Given that a traveling salesman must visit each of a number of cities, how should he choose his route in order to minimize the distance traveled? If the number of cities is small–three to ten, say–a computer can simply examine all the possible routes and calculate which one is shortest. But as the number of cities increases, the number of possible routes increases much faster. For a problem involving 100 cities, there are about 10^{155} possible routes. This number is larger than William A. Dembski's "universal probability bound" of 10^{150}, which can be roughly described as the outer limit of the number of potentially meaningful events in the history of the universe (Dembski, 1998a, 207-210). No computer could ever perform this number of operations, let alone examine this many routes.

Fogel's evolutionary algorithm to solve the 100-city traveling salesman problem starts with 100 routes chosen at random. Each route is represented by a sequence of numbers that describe the order in which the salesman visits each city. This sequence is the "gene" of that particular route. Fogel then generates 100 offspring routes which have partly the same genetic material as their parents and partly different material. He then examines the 200 parent and offspring routes to find which ones are shortest. He picks the 100 best ones, breeds another generation of 100 offspring routes, selects the 100 best in this second generation, and repeats the process for as many as 4000 generations. In this way, he arrives at solutions which, if not mathematically optimum, are very much better than the random guesses he started with. And he finds these within a very reasonable amount of computer time, because the breed-test-breed cycle is computationally straightforward and deals with only 200 routes at a time.

Evolutionary computing often finds solutions whose form surprises the designers of the algorithms. In an application of evolutionary computing to antenna design, engineers Edward Altshuler and Derek Linden commented that "the shape of the antenna was very unusual and could not likely have been obtained using other design methods." (Altshuler and Linden 1999, 223). Another design, which they called a "crooked wire genetic antenna," prompted them to comment that they found it "hard to believe that this odd-looking structure" met the design goals they established for the algorithms they used. But experimental measurements proved that it did (Altshuler and Linden 1999, 245).

While there are many forms that applied evolutionary computing can take, all successful applications have two characteristics in common: (1) they use some aspect of evolution's generational feedback mechanism to achieve incremental improvements in each generation with respect to some desiderata, and (2) human beings establish the desiderata that evolutionary computing is designed to achieve. With these characteristics in mind, we shall next turn to two more basic questions: what do computers really do in general, and to what extent do computers that run evolutionary algorithms use randomness as a part of their operation? The answers to these questions will put us in a position to compare natural evolution with evolutionary algorithms.

3. The Nature of Computation

Everyone who has heard an advertisement for digital cell phones knows something about the difference between digital and analog electronic systems: digital is better. Exactly why digital is better is perhaps less clear. It is true that digital electronic impulses are easier to transmit and harder to degrade than analog impulses. This is one practical reason that digital machines largely displaced the older analog computers, which were quite popular up to the 1950s. Nevertheless, there is an important sense in which all computers–digital, analog, slide rules, and abacuses–are “analogue” computers. I will revert to the archaic spelling to distinguish between the conventional sense of analog and the special sense of “analogue” which I will now develop.

The point is that everything the material object called a computer does is an analogue to non-material, and usually mathematical, concepts. For example, as I depress certain keys on my laptop computer in a certain sequence, certain pixels on its screen go from white to black, and certain locations in its memory systems undergo electrical and magnetic changes, some of which remain over time. Materially, that is all the computer does. But in correspondence with those material processes is a non-material entity called “What Does Evolutionary Computing Say About Intelligent Design?” This non-material entity is the essay I am writing. Computers are useful for recording, storing, transmitting, and reproducing physical representations of non-material things such as essays precisely because there is a close correspondence, or analogy, between the non-material things called letters and words on the one hand, and certain electrical patterns, screen displays, and ink patterns on printer paper, on the other hand. (I assume that the reader does not subscribe to an absolute materialism which denies the existence of any non-material entities. Such a position is so difficult to hold with any consistency that I will not consider it further.)

An analogy is the setting up of a comparison between a first thing and a second thing, in the hopes that information known about the first thing will also provide information about the second thing. The two things must bear some relation or correspondence to each other, although that correspondence is generally limited. For example, a common analogy used in physics is the similarity between water waves in a still pond and other kinds of waves such as sound waves. Both waves share a number of common characteristics: periodicity in time and space, a specific relationship between wave amplitude and distance traveled from the source, and so on. But there are also differences: water waves are basically an interfacial phenomenon, sound waves generally propagate through a volume of space, and so on.

Since the advent of mathematical physics in the seventeenth century, one overarching analogy has become so commonplace that it is rarely thought of any more as an analogy. I speak of the representation of all sorts of physical quantities with numbers and mathematical statements about those numbers. The physical world, it is needless to say, is made up of material things. Numbers are non-material things. Yet it is one of the grand mysteries of the universe why so many diverse material phenomena, from subatomic particles to the motion of galaxy clusters, can be analyzed, modeled, and in many cases predicted with the non-material entities of mathematics.

If science has displaced theology as queen of the intellectual world, surely mathematics has pride of place among the sciences as being the “most scientific” of sciences. Valid mathematical proofs, being exercises in logic, compel the much-desired universal consent which is the purported goal of modern scientific discourse. And the more mathematical a science is, the more credibility and prestige it tends to receive in the scientific community. The mathematical and logical sophistication of arguments for intelligent design made by such investigators as Dembski has played an important role in gaining their ideas such attention as they have received.

Yet any science which deals with “mathematics and”–mathematics and physics, mathematics and chemistry, mathematics and biology–must necessarily draw analogies between the physical entities which are its subjects in the physical world, and the non-material entities of the mathematical models it uses. Those analogies which are successful in representing the physical world with mathematics command respect simply because they work, often extraordinarily well. And in this pragmatic age, anything which works well commands respect.

Now that I have established that all mathematical sciences (other than pure mathematics itself) are based on the principle of analogy, I shall return to the specific way in which computers form analogues of the real world.

Digital computers are a special way to represent numbers with physical tokens such as the presence or absence of a certain charge in a memory chip. When Altshuler and Linden calculated the radiation pattern of their genetic antenna, they used two analogies. One analogy was the relationship between how the real antenna radiates real radio waves, and how certain equations known as Maxwell’s equations predict that these waves should be radiated. Although they have limitations, Maxwell’s equations do a nearly perfect job of predicting the behavior of antennas under the conditions that the engineers were dealing with. The reason that no one questions the engineers' use of a "mere analogy" to predict the antenna's behavior from its shape, is that such analogies have been working very well for more than a century. The other analogy that the engineers used was between the numerical data corresponding to Maxwell's equations for the particular antenna, and the physical configurations of the computer which represented that numerical data. Both types of analogies are so well-known and accepted that few ever give them a thought. They are, nevertheless, analogies. Mathematical models of the physical world are not the same as the parts of the physical world that are modelled, and the physical computer configurations corresponding to those models are technically distinct from both the physical reality and the mathematical model.

Now let us consider a different aspect of the design of the genetic antenna: the random variation and selection that gave rise to the improved antenna designs. In contrast to the use of Maxwell's equations to model the antenna's radiation characteristics, the variation and selection process that took place on the computer does not model exactly any known physical process. However, it does imitate in a general way the genetic sequence of biological events in nature. I claim that the similarity between the way that variation and selection was used to design the genetic antenna, and the way that variation and selection operates in biology, warrants the extension of the analogy between evolutionary computation and natural biological evolutionary processes. In particular, because the computational evolutionary process used a small amount of randomness within a highly structured framework designed by minds, I claim that this fact provides evidence that mind is necessarily involved in the natural biological process generally termed evolution.

4. Mind and Design

To give concreteness to the discussion, I will re-examine the example given above of the antenna designers who used a genetic algorithm to design a new type of antenna. During this discussion, I will take special care to define design.

Del Ratzsch favors the following definitions of design (noun and verb form), which I will adopt (Ratzsch 2001, 3):

i. a *pattern* is an abstract structure which correlates in special ways to mind, or is *mind correlative*.

ii. a *design* is a deliberately intended or produced pattern.

iii. to be *designed* is to exemplify a design.

It is obvious from the first part of Ratzsch's definition that the concept of mind must enter into any discussion of design which uses his definitions of the word. We attribute intention and deliberation only to beings with minds. If we accept Ratzsch's definition of design, the presence of a design implies the existence of a designer: a mind which originated the design.

And here we encounter the question of origin. Since most people have great difficulty in explaining where their original ideas come from, it is customary to terminate the causal chain that emerges from a mind at the mind itself. That is, when I say that the origin of Beethoven's Ninth Symphony was Beethoven, I mean that the idea of the symphony originated in Beethoven's mind. While there is a sense in which "there is nothing new under the sun," original works such as symphonies and genetic antennas are generally accepted to originate in the minds of their designers.

If we accept the classification of the genetic antenna as a design, and use Ratzsch's definition of design, we must conclude that the antenna was designed not by a computer, but by its human designers. In this way of thinking, the computer was simply a tool, albeit a very sophisticated one, in the hands of the human designers. This is true even though the designers admit to some surprise at the exact form of the antenna, which they could not have anticipated in advance of the actual evolutionary computation that produced it.

Let us summarize what we have developed thus far. For the purposes of this discussion, the genetic antenna's origin was in the minds of the designers, who embodied their ideas in certain non-material computer programs which in turn were transformed into material arrangements of matter in a physical computer, which in turn produced a physical plan embodying the non-material antenna design, from which a technician made a physical antenna which fulfills a certain purpose. This purpose was originally in the minds of the designers, and cannot be detected readily at the level of the computer program, except implicitly in the selection criteria programmed into the genetic algorithm.

The story just summarized comprises a description of information origin and flow. It will be instructive to trace this origin and flow in some detail in chronological order. At some point in the past, Altshuler and Linden "had an idea" about designing a genetic antenna. At first, in the nature of such ideas, its outlines were indistinct, few details were clear, and early in the design process no physical antenna could have been made because insufficient information was available.

The designers then wrote (and rewrote, no doubt) a set of computer algorithms that included a set of criteria or design goals, a process by which a design could be evaluated against these goals, and a genetic algorithm which could start with a given design, propagate and vary generations, test them against the goals, and select the best designs for future development in the same way. This procedure of writing the programs was highly interactive. At each step, the human designers tried something on the computer, corrected mistakes, observed the results, and modified the program until it began doing what they wanted it to do. In this sense, the writing of any computer program is an evolutionary process in which various tries are sorted through, modified, and improved until a satisfactory result is achieved. Very few if any programmers or computer scientists sit down and write an entire software package from scratch without a mistake, and have it run properly the first time.

Notice, however, that at the end of the process, when the genetic algorithm had produced a good antenna design, more information existed than before the process began. At the beginning of the process, there were only vague ideas in the designer's minds. At the end, there was a highly specific set of ideas corresponded to by a highly specific arrangement of matter both in the computer, on the set of plans, and in the genetic antenna itself. Now for the prize question: where did this information come from? What was its origin?

Questions about origin are really questions about causal chains. We have already traced the causal chain for the antenna itself, the physical object. The information which the antenna embodies has a rather peculiar history. From the viewpoint of the designers, the specific information about how the antenna should be formed gradually emerged over a series of many generations of genetic algorithms. In some sense, the final explicit or actual information at the end of the process was implicit or potentially in the system at the beginning. But at the start, it had neither actual physical existence (in the computer) nor non-material existence in the minds of the designers, except as vague hopes and intentions which the designers admit did not have a specific form.

One way to think of causal chains is to pose counterfactual situations in which one removes one or more factors which were actually present in the real causal chain, but which are absent in the counterfactual one. One then asks about the counterfactual chain, “Would the same outcome have occurred if this counterfactual chain were the case instead of the actual one?” If the answer is “No,” then the factor removed must be considered as a cause of the outcome, in some sense.

For example, if I get up in the morning and trip over my son’s toy wagon, I say that the wagon caused me to trip. I say this because if I imagine a counterfactual sequence of events which are exactly the same except for the absence of the wagon, it is highly likely that I would not trip, and the same outcome would not have occurred. Therefore, I am correct in saying that the wagon caused me to trip.

Now it is true that the computer, the technician, and even the wire used to make the genetic antenna are all “causes” in this counterfactual-removal sense. But common speech distinguishes a ranking or ordering of causes, generally favoring more significant and contingent causes over others. Minds, with their apparent free will and freedom from the mechanical and physical necessities of cause and effect, are especially liable to have the word “cause” attached to them. If an entity could have either caused an event or not, but chose to, we are especially likely to say that this entity was *the* cause.

Physical inevitabilities do not generally get described as causes. If a man starts a forest fire, we do not hear the reporters say, "A lighted cigarette carelessly thrown in the bushes caused the fire," unless the lighted cigarette is as far back as the causal chain can be traced. Obviously, the cigarette didn't throw itself away: some person did, and that person is generally regarded as the cause of the blaze.

While the human antenna designers were the cause of the genetic antenna in this common-sense meaning of the word, it is false to attribute to them directly every bit of information necessary to build the antenna. Instead, they used a tool called a computer to produce more information than they were able to produce without its aid. This information gradually came into being, guided at all times by the human designers and selected with the antenna's purpose in mind.

A reviewer of an earlier version of this essay urged upon the author the distinction between origin in the sense of actualization, and origin in the sense of ultimate ground of being. These different terms can be usefully applied to the case of the genetic antenna. Since the human designers, unaided by the computer, were unable to devise the antenna's exact form, but the computer did, it would be fair to say that the computer actualized the antenna. (Technically, the technician did, but if we suppose there is an automated antenna-bending rig attached to the computer, we eliminate this minor participation of humans in the final outcome.) But if we take origin in the sense of ultimate ground of being, then (to the extent humans are considered ultimate), the human designers were the origin of the antenna. They conceptualized it in a general sense, while the computer actualized it.

With these distinctions in mind, we next turn to the question of variation and natural selection in nature and how it compares with the computer-based genetic process we have discussed up to now.

5. Natural Selection: A Natural Genetic Algorithm?

Consider the well-documented process by which bacteria can evolve a resistance to a certain hostile chemical or antibiotic in their environment. This is an experiment which has been demonstrated to the satisfaction of all concerned, and even intelligent-design advocates such as Michael Behe admit that such “microevolutionary” processes occur. In what ways is it analogous to the process of designing an antenna with a genetic algorithm?

Only while the computer is actively producing generations of improved antennas does the computational process have even an approximate parallel in biological nature. This is because every known living species has existed since long before mankind began to be scientifically curious about the origin of life. No human was around to watch life begin. By contrast, we have an essentially complete record not only of the particular genetic antenna design under consideration, but of the development of software algorithms and computers in general. If we were to artificially restrict our knowledge about the genetic antenna’s origin so that it was comparable to what we know about biological origins, we would allow ourselves only to observe the last few generations of antenna designs being produced, and perhaps allow ourselves a small peek at scraps of the source code, comparable to the very incomplete nature of the fossil record. With these limitations in mind, let us see how many parallels we can find between the case of evolutionary computing and the case of evolving bacteria.

First, we will identify analogous physical structures and their purposes. Information is embodied in both the bacteria and in the computer. In bacteria, some information is contained in DNA, but by no means all of the information necessary to construct a bacterium from scratch. Even outside the bacteria's cell walls, yet more information is contained in the environment in which the hostile chemical resides. If we do not include the bacteria's environment, the conditions necessary for the kind of evolution we are looking for will be incomplete. Without the external pressure of the hostile chemical, the bacteria will mutate randomly, but will not evolve the new characteristic of being resistant to the chemical.

In the case of the genetic antenna, the desired antenna characteristics are embodied in a different part of the computer program than the part where the variation and propagation of the algorithm's generations take place. One part handles the mechanics of reproduction and variation and another part handles the selection criteria. So the analogue to the hostile chemical in the case of the bacteria, is the selection criteria built into the program by the human designers.

Another physical artifact which the bacteria and the antenna cases share is the presence and dissipative use of energy. A static, fixed-energy system cannot evolve, because nothing can change. A seldom-considered but universal characteristic of all life processes is the production of waste products, both waste energy and waste matter. Of course, the definition of "waste" may be philosophically problematic, but we use the common-sense definition that waste is material produced as a by-product of other processes more central to the continued existence of the organism, including reproduction itself. Interestingly, although it is theoretically possible to do computation with zero wasted energy, this is possible only if no memory storage elements are ever erased and rewritten. Since every practical computing machine must erase and reuse memory instead of just letting it pile up indefinitely, it has been proved that practical computing machines also require energy and must increase the net entropy of the universe, much as other heat engines do. This means that for either biological or computational evolution to occur, energy must be provided to the system. This fact has interesting implications related to the second law of thermodynamics and the direction of time, which we will not pursue. Nevertheless, it is another parallel that reinforces the fundamental similarity between biological evolution and computational evolution, in some sense.

The purpose of making analogies is to use knowledge developed in one domain to analyze and predict phenomena in a second analogous domain. In the case of evolutionary computing, we have examples of a type of evolution which are known exhaustively, from initial programming by human minds to the final outcome of the designed physical artifact. In particular, we know beyond the shadow of a doubt that the entire process, including the machine on which the program runs, originated in the minds of humans. There is not a single known case of computers spontaneously programming themselves with no human intervention whatsoever. Although computer programs may do things which surprise their programmers, and programs have been written which in turn write other programs, the human mind is the source of all the basic structures and constraints under which computer programs operate. As far as we know, if there were no human minds, there would be no computers and no programs.

What about the case of biological evolution? In the specific example of the bacteria which evolve a resistance to a hostile chemical, if one compares the information content in the bacteria's DNA after this process to the information content at the beginning of the process, one will presumably find that the DNA has changed. The information is different, but whether there is, strictly speaking, more or less information would be a difficult question to answer.

Claude Shannon's rudimentary measure of information in information theory is framed in terms of information as events and the probability of such events. If an event occurs with probability P , then the amount of information associated with that event is defined as $I = \log_2(1/P)$, where the unit of information is the bit. The difficulty in applying this measure of information to situations as complex as evolution (of either bacteria or antennas) is that the calculation of the probabilities involved quickly becomes unmanageable. If a certain artifact has never existed before, there is no basis of data with which to calculate the probabilities in advance of the event's occurrence, and therefore no way to arrive at a meaningful measure of information in the Shannon sense. Nevertheless, we feel intuitively that a sequence of DNA several million base pairs long contains more information than, say, a perfect crystal of silicon. And the DNA in a bacterium which, in addition to its other characteristics, can resist a hostile chemical, presumably has some feature that enables this resistance, a feature that was not present before the evolutionary change took place. This feature we will take to be additional information.

Another important fact from computer science may shed light on the biological case. A sequential-program computer in good working order, once it is informationally isolated from its environment, is a perfectly deterministic machine. That is, if one programs it a certain way and then allows it to run a program to completion, the program will always turn out the same way. Strange as it may seem, this includes the production of genetic antennas.

This principle of determinism seems at first to contradict the idea that we can do "random" mutations and variations on genetic algorithms. But no digital computer can produce a truly random outcome without a random input from outside its deterministic world. There are algorithms for generating pseudo-random sequences of numbers, but even these eventually repeat themselves. If one initializes such a random-number generation program the same way every time, it will always generate the same sequence of "random" numbers, which are, of course, not truly random. There are ways of producing more nearly random numbers, but they all involve some input from the non-mathematical world such as a signal from an analog random-noise generator. If such inputs are used, then the computer is no longer strictly deterministic and there is no way to say with assurance how a program using such physically random inputs will turn out.

We have arrived at an odd situation. On the one hand, we have an evolutionary computation which will always turn out the same way unless we intrude into the computer's perfect, mathematically deterministic world with a random input from the real world. Even if we do, the evolutionary computation still tends to produce suitable antennas, although they may differ in form slightly from one design attempt to the next if we use physically random inputs.

On the other hand, we have a real case of biological evolution of bacteria. We can consider this experiment as a biological "computation" that ends up producing more information than we had at the beginning. If we consider the bacteria's newly developed ability to resist the chemical a new piece of information, then we can say that there is more information contained in the bacteria's genetic material than before.

But is it not the case that the added genetic information was present in a potential form in the environment itself? We needed only to "run the computation" (that is, perform the experiment of letting the bacteria propagate over several generations) to make the potential information actual. In the case of the computational genetic process, it is a mathematical certainty that the final outcome of the computation resided implicitly in the initial conditions of the computer at the start of the process. While it is less than a mathematical certainty that the bacteria would evolve the way they did, one can take identical strains of bacteria and repeat the experiment with substantially the same outcomes. This may be as close to mathematical certainty as biology comes.

Let us summarize the similarities between the two evolutionary processes. They are:

1. Both processes produce a sequence of generations whose characteristics varied, and which were selected by an externally applied criterion which moved the evolutionary process in a certain direction.
2. Both processes require and dissipate energy to move the physical part of the process forward in time.
3. Both can be considered as computations.
4. Both processes ended with more explicit information present than when they began.
5. In both cases, the additional information present at the end of the computation was implicit in the experiment's initial conditions.

So much for the similarities. What do we know about the evolutionary computation of the antenna design that we do not know about the evolution of the bacteria? Several things:

1. Minds originated the entire plan for the antenna design, determined the selection criteria, and designed the software and hardware on which the evolutionary process took place.
2. Although random inputs played a role in the genetic antenna design, they were a small part of an otherwise highly structured process which issued in a product whose general characteristics were foreknown by the designers, although the specific details were not.

I will now return to Ratzsch's definition of design I quoted earlier, and ask the question: Do the chemical-resistant bacteria at the end of the biological experiment qualify as designed? They are designed to the extent that their newly acquired resistance was a property that the human designer of the experiment wished them to have. In principle, this experiment is no different than breeding cows or roses for certain desirable characteristics. Darwin himself begins *The Origin of Species* with a long discussion of breeding practices. The critical difference between the bacteria experiment and the antenna experiment is that bacteria were already a going concern when the scientist came along to experiment on them. The production of the genetic antenna was a much more *ab initio* process, and minds were involved at the start as well as during the evolution itself.

Everyone will agree that the chemical resistance the bacteria acquired is a "designed" characteristic, produced indirectly by the breeding process which the experimenter created the conditions for. But there is much disagreement over the question of whether bacteria themselves, or any other living entity, embody design in the sense in which Ratzsch and I use it.

The most important conclusion we can draw from the analogy between evolutionary computation and natural evolutionary processes such as the one I have described, is as follows. In absolutely every case of evolutionary computation, without exception, minds are the origin of the vast majority of the features of the evolutionary process, the selection criteria, and the computing environment in which the evolutionary process takes place. Although a certain amount of randomness is involved in evolutionary computing, it is small compared to the exacting precision of the computations and evaluations that must take place for successful evolutionary computation to occur.

I leave it to the reader to draw the analogous conclusion for the case of biology. If we have thousands of cases of exhaustively known computational evolution going on around us on computers, and we know for a fact that minds were involved in the origin of every single one, the conclusion that mind was involved in the origin of naturally occuring species is an obvious one. This analogy says nothing, unfortunately, about the <u>means</u> which mind might have used to originate species. Since, under this hypothesis, we ourselves are presumably one of the species which mind originated, we are somewhat in the position of yet another genetic algorithm on a computer, an intelligent one which can speculate about its own origin. In an earlier version of this essay, I attempted to develop this idea, but it requires that we assume that a computer program has the attributes of self-consciousness and mind. I find insufficient evidence to warrant this assumption, so I will simply say that the analogy between evolutionary computation and biology breaks down at this point. All analogies have limits beyond which they fail, and this is one limitation of the present analogy.

6. Some Objections

I would like to anticipate some objections to my argument that evolutionary computing provides us with analogical evidence that mind is involved in biological evolution.

One objection to this conclusion is that mind and intelligence are products of evolution, and so evolution is the origin of mind, not the other way around. This point of view is represented by, among others, David Fogel, a leader in the field of evolutionary computation. Fogel claims that intelligence is "the capability of a system to adapt its behavior to meet its goals in a range of environments." (Fogel, 1995, 24) This rather broad definition ascribes some degree of intelligence to all living things and many artificial systems, even those as simple as a home heating system regulated by a thermostat. With such a broad definition, Fogel has no problem in saying that evolutionary algorithms are intelligent.

What about the process of evolution itself? According to Fogel, evolution, which he defines in a fairly standard neo-Darwinist way, is highly intelligent. But going beyond that, Fogel puts all intelligent processes under the umbrella of evolution when he says, "Evolution serves as a unifying description of all intelligent processes." He views all forms of intelligence through a Darwinian lens: " . . . every intelligent system adopts a functionally equivalent process of reproduction, mutation, competition, and selection. . . ." (Fogel, 1995, 249) According to Fogel, evolution is more than simply a biological process. It is the overarching pattern of all intelligence in humans, other animals, and machines. Therefore, according to Fogel, evolution is the origin of mind, intelligence, and all living things.

Just as there are some who will refuse to accept the implications of an analogy simply because it is an analogy, there are others who insist on definitions which allow them to maintain their point of view. Materialists and others who believe that mind is an epiphenomenon of the physical brain will not be convinced by my arguments about mind being the origin of evolutionary processes in nature as well as in computers. If one's postulates about the universe exclude the possibility that mind, of a nature not yet comprehended by us, could be responsible for features of the observable universe, no evidence to the contrary will be convincing.

A different but related objection to my thesis is that while it may be the case that mind has something to do with biological evolution, it makes no difference to science whether this is the case or not. William James, in *Principles of Psychology*, argued along similar lines in his discussion of whether the concept of the soul was useful to scientific psychology (James, 1890, 325-332). He concluded that while the then-current state of psychological knowledge was incapable of disproving the existence of the soul, the hypothesis was scientifically fruitless. According to James, if we assume that there is such a thing as a soul, it provides no predictions or models which can be empirically verified.

We are in a similar situation with regard to the hypothesis that mind is associated with biological evolution. One of the postulates of neo-Darwinian evolution is that evolution has no purpose. Most Darwinians agree that it is a random, meaningless, pointless process with no higher motive or guiding principles. This postulate is stated in the face of the incredible order, hierarchy, and system of the biological world. It may appear very improbable that such an organized world is the product of randomness, but the postulate of non-purposiveness requires it.

If instead, we adopt the hypothesis that mind directed and directs biological development by some as-yet-undiscovered means, we have not solved the problem of the incredible order and complexity of biology. We have simply moved it back one step to the question of the nature of the mind which brought it about. Unless we possess independent data about the nature of such a mind, we are basically in the same position as those who believe biological evolution has no purpose. So the bare statement that mind is behind evolution is no threat to the postulate of randomness behind evolution, until we begin to attribute characteristics to the mind we postulate. And unless such characteristics can be determined within the empirical framework of science, the hypothesis of mind will turn out to be as scientifically barren as William James regarded the hypothesis of the soul in psychology.

In fact, most biologists already believe in a practical way that some sort of intelligence, if not mind, is behind biology. Otherwise, they would not seek the inner mechanisms that form the content of so much modern research. These mechanisms fit intellectual patterns which biologists first hypothesize, then (with any luck) discover to be the case by means of successful experiments. While it is true that when asked, most biologists attribute the source of these structures to evolution (meaning, presumably, randomness), they could just as easily attribute them to mind. The fact that they do not, is not a measure of their scientific prowess or lack of it. Instead, it expresses their philosophical position of metaphysical naturalism. The relatively few biologists such as Michael Behe who attribute the exact same scientific features of biology to the mind of God, do not do science any differently. They simply maintain a different philosophical and religious orientation.

It remains to be seen whether the hypothesis of mind behind biological evolution will prove to be scientifically fruitful. In my opinion, it already has proved to be so. The roots of modern science lie deep in the Christian worldview of a God whose intellect orders the universe. That belief in a mind-pattern that underlies the visible world is the only thing that keeps science going. To the extent that scientists *genuinely* begin to believe that randomness is supreme, they will cease to seek pattern and order in the world, and science will slow down and eventually cease. The fact that this has not yet happened tells us that belief in an ordering mind behind biology has not yet vanished from the face of the earth. Nevertheless, simply for the purposes of philosophical solidarity, most biologists give at least lip service to the idea of randomness as the supreme driver of biological evolution, while believing heartily in their own mind-patterns concerning the natural world that turn out to be verified by experiment.

7. Conclusions

It is often the case that an analogy in science provides a prediction long before the mechanism behind the prediction is elucidated. Planck's hypothesis of the quantum arose from his discovery that it resolved certain difficulties in radiation physics, but neither Planck nor anyone else understood at first why or how energy might have to be exchanged in discrete quantities rather than continuously. The sciences of mind and consciousness are yet in an early stage of development, and it may be some time before we can develop precise enough measures to reach scientifically valid conclusions about the involvement of mind in the origin of particular artifacts. But my argument by analogy from evolutionary computation is that the fantastic complexity and variety of living beings, as well as the processes by which they propagate, preserve, and modify genetic information, must have originated in mind. The evidence I cite is the fact that every known case of computational evolution originated in the minds of humans, and uses so-called randomness as an essential but relatively minor input.

To believe that the process of natural evolution *as it can be observed today in the laboratory*, is sufficient to account for the origin and variety of species, is like believing that the random input to the genetic algorithm on a computer is wholly responsible for not only the observed modifications and improvements from generation to generation, but also for the design of the entire genetic algorithm and the computer itself. The analogy I put forward does not provide details about the mechanism by which mind may have influenced the origin and variety of species. But it does imply strongly that the present hypotheses, based on metaphysical naturalism and random processes alone, are inadequate.

8. Acknowledgments

The author thanks William A. Dembski and Calvin College's Seminars in Christian Scholarship Series, which provided support and encouragement for this research during the summer of 2000, and the anonymous reviewers of an earlier version of this essay for their helpful criticisms.

Endnotes

[1] A recent sampling of books on genetic algorithms at a large research library (the University of Texas at Austin) turned up 48 titles and 29 conference proceedings.

References

Altshuler, Edward E., and Derek S. Linden. 1999. "Design of Wire Antennas Using Genetic Algorithms." In *Electromagnetic Optimization by Genetic Algorithms* , eds. Yahya Rahmat-Samii and Eric Michielssen, 211-248. New York: J. Wiley & Sons.

Dembski, William A. 1998a. *The Design Inference* . New York: Cambridge Univ. Press.

Fogel, David B. 1995. *Evolutionary Computation: Toward a New Philosophy of Machine Intelligence.* Piscataway, NJ: IEEE Press.

Fogel, David B. 2000. "What is evolutionary computation?" *IEEE Spectrum* 37 (Feb.):26-32.

Holland, John H. 1975. *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: Univ. of Michigan Press.

Holland, John H. 1992. " Genetic Algorithms." *Scientific American* 267 (July):66-72.

James, William. 1890. *Principles of Psychology*, vol. 1 (reprinted 1981). Cambridge, MA: Harvard University Press.

Ratzsch, Del. 2001. *Nature, Design, and Science: The Status of Design in Science*. Albany, NY: State University of New York Press.