Hillis, D. W., Intelligence as an Emergent Behavior; or, The Songs of Eden, p. 175-189, *Dædalus*, Journal of the American Academy of Arts and Sciences, special issue on Artificial Intelligence, Winter 1988

Intelligence as an Emergent Behavior; or, The Songs of Eden

W. Daniel Hillis

Sometimes a system with many simple components will exhibit a behavior of the whole that seems more organised than the behavior of the individual parts. Consider the intricate structure of a snowflake. Symmetric shapes within the crystals of ice repeat in threes and sixes, with patterns recurring from place to place and within themselves at different scales. The shapes formed by the ice are consequences of the local rules of interaction that govern the molecules of water, although the connection between the shapes and the rules is far from obvious. After all, these are the same rules of interaction that cause water to suddenly turn to steam at its boiling point and cause whirlpools to form in a stream. The rules that govern the forces between water molecules seem much simpler than crystals or whirlpools or boiling points, yet all of these complex phenomena are somehow consequences of those rules. Such phenomena are called *emergent behaviors* of the system.

It would be very convenient if intelligence were an emergent behavior of randomly connected neurons in the same sense that snowflakes and whirlpools are emergent behaviors of water molecules. It might then be possible to build a thinking machine by simply hooking together a sufficiently large network of artificial neurons. The notion of emergence would suggest that such a network, once it reached some critical mass, would spontaneously begin to think.

This is a seductive idea because it allows for theh possibility of constructing intelligence without first understanding it. Understanding intelligence is difficult and probably a long way off, so the possibility that it might spontaneously emerge from the interactions of a large collection of simple parts has considerable appeal to the would-be builder of thinking machines. Unfortunately, that idea does not suggest a practical approach to construction. The concept of emergence in itself offers neither guidance on how to construct such a system nor insight into why it would work.

Ironically, the apparent inscrutability of the idea of intelligence as an emergent behavior accounts for much of its continuing popularity. Emergence offers a way to believe in physical causality while simultaneously maintaining the impossibility of a reductionist explanation of thought. For those who fear mechanistic explanations of the human mind, our ignorance of how local interactions produce emergent behavior offers a reassuring fog in which to hide the soul.

There has been a recent renewal of interest in emergent behavior in the form of simulated neural networks and connectionist models, spin glasses and cellular automata, and evolutionary models. Each of these is a model of some real system. For neural networks and connectionist models, the system being modeled is a collection of biological neurons, such as the brain; for spin glasses it is molecular crystals. Cellular automata and evolutionary models are based on the ontogenesis and phylogenesis of living organisms. In all of these cases, both the model and the system being modeled produce dramatic examples of emergent behavior.

Most of these models are not new, but interest in them is being stirred because of a combination of new insights and new tools. The insights come primarily from a branch of physics called dynamical systems theory. The tools come from the development of new types of computing devices. Just as we thought of intelligence in terms of servomechanism in the 1950s, and in terms of sequential computers in the sixties and seventies, we are now beginning to think in terms of parallel computers, in which tens of thousands of processors work together. This is not a deep philosophical shift, but it is of great practical importance, since it is now possible to study large emergent systems experimentally.

Inevitably, antireductionists interpret such progress as a schism between symbolic rationalists, who oppose them, and gestaltists, who support them. I have often been asked which “side” I am on. Not being a philosopher, my inclination is to focus on the practical aspects of this question: How would we go about constructing an emergent intelligence? What information would we need to know in order to succeed? How can this information be determined by experiment?

The emergent system that I can most easily imagine would be an implementation of symbolic thought rather than a refutation of it. Symbolic thought would be an emergent property of the system. The point of view is best explained by the following parable about the origin of human intelligence. As far as I know, this parable of human evolution is consistent with the available evidence (as are many others), but because it is chosen to illustrate a point, it should be read as a story, not as a theory. It is different from most accepted theories of human development in that it presents features that are measurable in the archaeological records—such as increased brain size, food sharing and neoteny—as consequences rather than causes of intelligence.

Once upon a time, about two and a half million years ago, there lived a race of apes that walked upright. In terms of intellect and habit they were similar to modern chimpanzees. The young apes, like young apes today, had a tendency to mimic the actions of others. In particular, they had a tendency to imitate sounds. If one ape shrieked “ooh, eeh, eeh,” another would repeat “ooh, eeh, eeh.” Some sequences of sounds, or "songs", were more likely to be mimicked than others.

Let us ignore the evolution of the apes for the moment and consider the evolution of the songs. Since the songs were replicated by the apes, and since they sometimes died away and were occasionally combined with others, we may loosely consider them (very loosely) a form of life. They survived, bred, competed with one another, and evolved according to their own criterion of fitness. If a song contained a particularly catchy phrase that caused it to be repeated often, then that phrase was likely to be incorporated into other songs. Only songs that had a strong tendency to be repeated survived.

The survival of a song was only indirectly related to the survival of the apes; it was more directly affected by the survival of other songs. Since the apes were a limited resource, the songs had to compete with one another for a chance to be sung. One successful competition strategy was for a song to specialize; that is, for it to find a particular niche in which it was apt to be repeated. Songs that fit particularly well with specific moods or activities of apes had a special survival value for this reason. (I do not know why some songs fit well with particular moods, but since it is true for me, I do not find it hard to believe that it was true for my ancestors.)

Before songs began to specialize they were of no particular value to the apes. In a biological sense the songs were parasites, taking advantage of the apes' tendency to imitate. As songs became specialized, however, it became advantageous for apes to pay attention to the songs of others and to differentiate between them. By listening to songs, a clever ape could gain useful information. For example, an ape could infer that another ape had found food or that it was likely to attack. Once the apes began to take advantage of the songs, a symbiotic relationship developed: songs enhanced their own survival by conveying useful information to apes; apes enhanced their own survival by improving their capacity to remember, replicate, and understand songs. Thus the blind forces of evolution created a partnership between the songs and the apes that thrived on the basis of mutual self-interest. Eventually this partnership evolved into one of the world's most successful symbionts: the human race.

Unfortunately, songs do not leave fossils, so unless some natural process has left a phonographic trace, we may never know if the preceding story describes what really happened. But if the story is true, the apes and the songs became the two components of human intelligence. The songs evolved into the knowledge, mores, and mechanisms of thought that together are the symbolic portion of human intelligence. The apes became apes with bigger brains, perhaps optimized for late maturity so that they could learn more songs. *Homo sapiens* is a cooperative combination of the two.

It is not unusual in nature for two species to live together so interdependently that they appear to be a single organism. Lichens are symbionts of a fungus and an alga that live so closely intertwined that they can only be separated under a microscope. Bean plants need living bacteria in their roots to fix the nitrogen they extract from the soil, and in return the bacteria need nutrients from the bean plants. Even the single-celled *Paramecium bursarra* uses green algae living inside itself to synthesize food.

Another example of two entirely different forms of "life" that form a symbiosis may be even closer to the example of the songs and the apes. In *The Origins of Life*, Freeman Dyson suggests that biological life is a symbiotic combination of two different self-reproducing entities with very different forms of replication.[1] Dyson suggests that life originated in two stages. While most theories of the origin of life start with nucleotides replicating in some "primeval soup", Dyson's theory starts with metabolizing drops of oil.

In the beginning these hypothetical replicating oil drops had no genetic material, but were self-perpetuating chemical systems that absorbed raw materials from their surroundings. When a drop reached a certain size it would split; about half of its constituents would go to each part. Such drops evolved efficient metabolic systems even though their rules of replication were very different from the Mendelian rules of modern life. Once the oil drops became good at metabolizing, they were infected by another form of replicators that, like the songs, had no metabolism of their own. These were parasitic molecules of DNA; like modern viruses, they took advantage of the existing machinery of the host cells to reproduce. The metabolizers and the DNA eventually coevolved into the mutually beneficial symbiosis that we know today as life.

This two-part theory of life is not conceptually far from the two-part story of intelligence. Both suggest that a preexisting homoestatic mechanism was infected by an opportunistic parasite. The two parts reproduced according to different sets of rules, but were able to coevolve so successfully that the resulting symbiont appears to be a single entity. Viewed in this light, choosing between emergence and symbolic computation in the study of intelligence is like choosing between metabolism and genetic replication in the study of life. Just as the metabolic system provides a substrate in which the genetic system can work, so an emergent system may provide a substrate in which the symbolic system can operate.

Currently the metabolic system of life is far too complex for us to fully understand or reproduce it. By comparison the Mendelian rules of genetic replication are almost trivial, and it is possible to study them as a system unto themselves without worrying about the details of the metabolism that supports them. In the same sense, it seems likely that symbolic thought can be fruitfully studied and perhaps even recreated without worrying about the details of the emergent system that supports it. So far this has been the dominant approach in AI and the approach that has yielded the most progress.

The other approach is to build a model of the emergent substrate of intelligence. This artificial substrate for thought would not need to mimic in detail the mechanisms of the biological system, but it would need to exhibit those emergent properties that are necessary to support the operations of thought.

What is the minimum that we would need to understand in order to construct such a system? For one thing, we would need to know how big a system to build. Information theory suggests that the appropriate unit of measure is the number of binary digits, or bits, required to store the information. How many bits are required to store the acquired portion of human knowledge of a typical human? We need to know an approximate answer in order to construct an emergent intelligence with humanlike performance. Currently the amount of acquired information stored by an average human brain is not known to within even two orders of magnitude, but it can in principle be determined by experiment. There are at least three ways to estimate the storage requirements for emergent intelligence.

One way would be through an understanding of the physical mechanisms of memory in the human brain. If information is stored primarily by modifications of synapses, then it would be possible to measure the information-storage capacity of the brain by counting the number of synapses. Elsewhere in this issue of *Dædalus*, Jacob T. Schwartz estimates that the brain contains roughly 10^{15} synapses. But even knowing the exact amount of physical storage in the brain would not completely answer the question of storage requirement, since much potential storage capacity might be unused or used inefficiently. But at least this method can help establish an upper bound on the requirements.

A second method for estimating the storage requirements for emergent intelligence is to measure the information in symbolic knowledge by some form of statistical sampling. For instance, it is possible to estimate the size of an individual's vocabulary by testing him or her on words randomly sampled from a dictionary. The fraction of test words known by the individual is a good indication of the fraction of words that he or she knows in the complete dictionary. The estimated vocabulary size is the test fraction multiplied by the number of words in the dictionary. Such an experiment depends on having a predetermined body of knowledge against which to measure. For example, it would be possible to estimate how many facts in the *Encyclopaedia Britannica* were known by a given individual, but this would give no measure of facts known by the individual but not contained in the encyclopedia. This method is useful only in establishing a lower bound.

A related experiment is the game of twenty questions, in which one player identifies an object chosen by another by asking a series of twenty yes-or-no questions. Since each answer provides no more than a single bit of information, and since skillful players generally need to ask almost all of the twenty questions to correctly identify the chosen object, we can estimate that the number of allowable choices known in common by the two players is on the order of 2^{20}, or about one million. Of course, this measure is inaccurate because the questions are not perfect and the choices of objects are not random. It is possible that a refined version of the game could be developed and used to provide another lower bound.

A third approach to gauging the human brain's storage requirements for information in the symbolic portion of knowledge is to estimate the average rate of information acquisition and to calculate the amount that would accumulate over time. For example, experiments on memorizing random sequences of syllables indicate that the maximum rate of memorization of this type of knowledge is about one "chunk" per second. A chunk, in this context, can be safely assumed to contain less than 100 bits of information, so the results suggest that the maximum rate at which a human is able to commit information to long term memory is significantly less than 100 bits per second.[2] If this is true, a twenty-year-old human learning at the maximum rate for sixteen hours a day (and never forgetting) would know less than 50 billion bits of information. I find this number surprisingly small.

A difficulty with this estimate of the rate of acquisition is that the experiment measures only information coming through one sensory channel under one particular set of circumstances. The visual system sends more than a million times this rate of information to the optic nerve, and it is conceivable that all of this information is committed to memory. If it turns out that images are stored directly, it will be necessary to significantly increase the 100-bit-per-second limit, but there is no current evidence that this is the case. In experiments measuring the ability of exceptional individuals to store eidetic (i.e., extraordinarily accurate and vivid) images of random-dot stereograms, the subjects are given about five minutes to memorize an image formed in a square array of 100 x 100 dots. Memorizing only a few hundred bits is probably sufficient to pass the test.

I am aware of no evidence that more than a few bits per second of any type of information can be committed to long-term memory. Even if we accept reports of extraordinary feats of memory (such as those of Luria's showman in *Mind of the Mnemonist*[3]) at face value, the average rate of commitment to memory never seems to exceed a few bits per second. Even if we knew the maximum rate of memorization exactly, the rate averaged over a lifetime would probably be very much less—but knowing the maximum rate would establish an upper bound on the requirements of storage.

The sketchy data cited above suggests that an intelligent machine would require 10^9 bits of storage, plus or minus two orders of magnitude. This assumes that the information is encoded in such a way that it requires a minimum amount of storage; for the purpose of processing information, this would probably not be the most practical representation. As a would-be builder of thinking machines, I find this number encouragingly small, since it is well within the range of current electronic computers. As a human with an ego, I find it distressing: I do not like to think that my entire lifetime of memories could be placed on a reel of magnetic tape. It is to be hoped that experimental evidence will clear this up one way or another.

There are a few subtleties in the question of storage requirements that involve defining the quantity of information in a way that is independent of its representation. Information theory provides a precise way of measuring information in terms of bits, but it requires a measure of the probabilities over the ensemble of possible states. That is, it requires assigning an a priori probability to each possible set of knowledge, which is the role of inherited intelligence. Inherited intelligence provides a framework in which the knowledge of acquired intelligence can be interpreted. Inherited intelligence defines what is knowable; acquired intelligence determines what of the knowable is known.

Another potential difficulty is how to count the storage of information that can be deduced from other data. In the strict information-theoretical sense, data that can be inferred from other data add no information at all. An accurate measure would have to take into account the possibility that knowledge is inconsistent, and that only limited inferences are actually made. These are the kinds of issues currently being studied on the symbolic side of the field of artificial intelligence.

One issue that does not need to be resolved to measure storage capacity is localized versus distributed representation—that is, whether each piece of information is stored in a specific place or spread “holographically” over a large area. Knowing what types of representation are used in what parts of the human brain is of considerable scientific interest, but it does not have a profound impact on the amount of storage in the system or on our ability to measure it. Nontechnical commentators have a tendency to attribute almost mystical qualities to distributed storage mechanisms such as those used in creating holograms and neural networks, but the limitations on the capacities of these storage mechanisms are well understood.

When a holographic plate is cut in two, each half contains a slightly degraded version of the entire image. Distributed representations with properties similar to holograms are often used within conventional digital computers, and they are invisible to most users except in the system's capacity to tolerate errors. The error-correcting memory system used in most computers is a good example. The system is composed of many physically separate memory chips, but any single chip can be removed without losing any data. This is because the data are not stored in any one place, but in a distributed, nonlocal representation across all of the units. In spite of this "holographic" representation, the information storage capacity of the system is no greater than it would be with a conventional representation, in which each piece of data is stored in a single chip. In fact, it is slightly less. This is typical of distributed representations.

Storage capacity offers one measure of the requirements of a humanlike emergent intelligence. Another measure is the required rate of computation. Here there is no agreed-upon metric, and it is particularly difficulty to define a unit of measure that is completely independent of representation. The measure suggested below is simple and important, if not sufficient.

Given an efficiently stored representation of human knowledge, what rate of access to that storage (in bits per second) is required to achieve humanlike performance? Here, *efficiently stored representation* means any representation requiring only a multiplicative constant of storage over the number of bits of information. This is a mathematical restriction that eliminates, for example, any representation that stores a precomputed answer to every question. Such a restriction does limit the range of possible representations, but it allows most representations that we would regard as reasonable. In particular, it allows both distributed and local representations.

The question of the memory bandwidth required for humanlike performance is accessible by experiment through approaches similar to those outlined for the question of storage capacity. If the time required for a primitive operation of human memory is limited by the firing time of a neuron, then the ratio of this "cycle time" to the total number of bits indicates what fraction of the memory is accessed simultaneously. This gives an indication of whether the brain is a parallel or serial device. In a serial device, data items are operated on sequentially, one at a time. In a parallel device, all data are operated on concurrently. Both serial and parallel behaviors are exhibited by the brain, but there is a question as to which model best describes the way that it reasons and accesses knowledge. Informed opinions differ greatly in this matter, but the bulk of the quantitative evidence favors serial computation. Memory retrieval times for items in lists, for example, depend on the position and the number of items in the list. Except for sensory processing, most successful artificial intelligence programs have been based on serial models of computation, although this may be a distortion caused by the common availability of serial machines.

My own guess is that the reaction-time experiments are misleading and that human-level performance will require that large fractions of knowledge be accessed several times per second. Given a representation of acquired intelligence with a realistic representation efficiency of 10 percent, the 10^9 bits of memory mentioned earlier would require a memory bandwidth of about 10^{11} bits per second. This bandwidth seems physiologically plausible, since it corresponds to about a bit per second per neuron in the cerebral cortex.

By way of comparison, the memory bandwidth of a conventional sequential computer is in the range of 10^6 to 10^8 bits per second. This is less than 0.1 percent of the imagined requirement. For parallel computers the bandwidth is considerably higher. For example, a 65,536-processor Connection Machine can access its memory at approximately 10^{11} bits per second.[4] It is not entirely coincidence that this fits well with the estimate above.

Another important question is, What sensory-motor functions are necessary to sustain symbolic intelligence? An ape is a complex sensory-motor machine, and it is possible that much of this complexity is necessary to sustain intelligence. Large portions of the brain seem to be devoted to visual, auditory and motor processing, and it is unknown how much of this machinery is needed for thought. A person who is blind and deaf or totally paralyzed can undoubtedly be intelligent, but this does not prove that the portion of the brain devoted to these functions is unnecessary for thought. It may be, for example, that a blind person takes advantage of the visual processing apparatus of the brain for spatial reasoning.

As we begin to understand more of the functional architecture of the brain, it should be possible to identify certain functions as being unnecessary for thought by studying patients whose cognitive abilities are unaffected by locally confined damage to the brain. For example, binocular stereo fusion is known to take place in a specific area of the cortex near the back of the head. Patients with damage to this area of the cortex have visual handicaps but show no obvious impairment in their ability to think. This is a simple example, and the conclusion is not surprising, but it should be possible by such experiments to establish that many sensory-motor functions are unnecessary. One can imagine metaphorically whittling away at the brain until it is reduced to its essential core. Of course, it is not quite this simple. Accidental damage rarely incapacitates a single area of the brain completely and exclusively. Also, it may be difficult to eliminate one function at a time because one mental capacity may compensate for the lack of another.

It may be more productive to assume that all sensory-motor apparatus is unnecessary until proven useful for thought, but this is contrary to the usual point of view. Our current understanding of the phylogenetic development of the nervous system suggests a point of view in which intelligence is an elaborate refinement of the connection between input and output. This is reinforced by the experimental convenience of studying simple nervous systems, or of studying complicated nervous systems by concentrating on those portions most directly related to input and output. By necessity, most everything we know about the function of the nervous system comes from experiments on those portions that are closely related to sensory inputs or motor outputs. It would not be surprising to learn that we have overestimated the importance of these functions to intelligent thought.

Sensory-motor functions are clearly important for the application of intelligence and for its evolution, but these issues are separate from whether sensory-motor functions are necessary for thought to exist. Intelligence would not be of much use without an elaborate system of sensory apparatus to measure the environment and an elaborate system of motor apparatus to change it, nor would it have been likely to evolve. But much more apparatus is probably necessary to exercise and evolve intelligence than to sustain it. One can believe in the necessity of the opposable thumb for the development of intelligence without doubting a human capacity for thumbless thought. It is quite possible that even the meager sensory-motor capabilities that we currently know how to create artificially would be sufficient for the fundamental operation of emergent intelligence.

Although questions of capacity and scope are necessary in defining the magnitude of the task of constructing an emergent intelligence, the key question is one of understanding. While it is possible that we will be able to recreate the emergent substrate of intelligence without fully understanding the details of how it works, it seems likely that we would at least need to understand some of its principles. There are at least three paths by which such understanding could be achieved. One is to study the properties of specific emergent systems—to build a theory of their capabilities and limitations. This kind of experimental study is currently being conducted on several classes of promising man-made systems, including neural networks, spin glasses, cellular automata, evolutionary systems, and adaptive automata. Another possible path to understanding is the study of biological systems, which are our only real examples of intelligence and our only examples of an emergent system that has produced intelligence. The disciplines that have so far provided the most useful information of this type have been neurophysiology, cognitive psychology and evolutionary biology. A third path would be a theoretical understanding of the requirements of intelligence or of the phenomena of emergence. Relevant examples are theories of logic and computability, linguistics, and dynamical systems theory. Anyone who looks to emergent systems as a way of defending human thought from the scrutiny of science is likely to be disappointed.

One cannot conclude, however, that a reductionist understanding is necessary for the creation of intelligence. Even a little understanding could go a long way toward the construction of an emergent system. A good example of this is how cellular automata have been used to simulate the emergent behavior of fluids. The whirlpools that form as fluid flows past a barrier are not well understood analytically, yet they are of great practical importance in the design of boats and airplanes. Equations that describe the flow of a fluid have been known for almost a century, but except for a few simple cases they cannot be solved. In practice the flow is generally analyzed by simulation. The most common method of simulation is the numerical solution of continuous equations.

On a highly parallel computer it is possible to simulate fluids with even less understanding of the system by simulating billions of colliding particles that reproduce emergent phenomena such as vortices. Calculating the detailed molecular interactions of so many particles would be extremely difficult, but a few simple aspects of the system, such as conservations of energy and particle number, are sufficient to reproduce the large-scale behavior. A system of simplified particles that obey these two laws but are otherwise unrealistic can reproduce the same emergent phenomena as reality. For example, it is possible to use particles of unit mass that move only at unit speed along a hexagonal lattice, colliding according to the rules of billiard balls.[5] Experiments show that this model produces laminar flow, vortex streams, and even turbulence that is indistinguishable from the behavior of real fluids. Although the detailed rules of interaction are very different from the interactions of real molecules, the emergent phenomena are the same. The emergent phenomena can be created without understanding the details of the forces between the molecules or the equations that describe the flow of the fluid.

The recreation of intricate patterns of ebbs and flows within a fluid demonstrates that it is possible to produce a phenomenon without fully understanding it. But the model was constructed by physicists who knew a lot about fluids. That knowledge helped to determine which features of the physical system were important to implement and which were not.

Physics is an unusually exact science. Perhaps a better example of an emergent system that we can simulate with only a limited understanding is evolutionary biology. We understand, in a weak sense, how creatures with Mendelian patterns of inheritance and different propensities for survival can evolve toward better fitness in their environments. In certain simple situations we can even write down equations that describe how quickly this adaptation will take place.[6] But there are many gaps in our understanding of the processes of evolution. We can explain why flying animals have light bones in terms of natural selection, but we cannot explain why certain animals have evolved flight while others have not. We have some qualitative understanding of the forces that cause evolutionary change, but (except in the simplest cases) we cannot explain the rate or even the direction of change.

In spite of these limitations, our understanding is sufficient to write programs of simulated evolution that show interesting emergent behaviors. For example, I have recently been using an evolutionary simulation to evolve programs to sort numbers. In this system, the genetic material of each simulated individual is interpreted as a program specifying a pattern of comparisons and exchanges. The probability of an individual survival in the system is dependent on the efficiency and accuracy of this program in sorting numbers. Surviving individuals produce offspring by sexual combination of their genetic material with occasional random mutation. After tens of thousands of generations, a population of hundreds of thousands of such individuals will evolve very efficient programs for sorting. Although I wrote the simulation that produced these sorting programs, I do not understand in detail how they were produced or how they work. If the simulation had not produced working programs, I would have had very little idea about how to fix it.

The fluid flow and simulated evolution examples suggest that it is possible to make a great deal of use of a small amount of understanding. The emergent behaviors exhibited by these systems are a consequence of the simple underlying rules defined by the program. Although the systems succeed in producing the desired results, their detailed behaviors are beyond our ability to analyze and predict. One can imagine that if a similar process produced a system of emergent intelligence, we would have a similar lack of understanding about how it worked.

My own guess is that such an emergent system would not be an intelligent system itself, but rather the metabolic substrate on which intelligence might grow. In terms of the apes and the songs, the emergent portion of the system would play the role of the ape, or at least that part of the ape that hosts the songs. This artificial mind would need to be inoculated with human knowledge. I imagine this process to be not so different from teaching a child. This would be a tricky and uncertain procedure because, like a child, this emergent mind would presumably be susceptible to bad ideas as well as good. The result would be not so much an artificial intelligence, but rather a human intelligence sustained within an artificial mind.

Of course, I understand that this is just a dream, and I will admit that I am propelled more by hope than by the probability of success. But if this artificial mind can sustain itself and grow of its own accord, then for the first time human thought will live free of bones and flesh, giving this child of mind an earthly immortality denied to us.

ENDNOTES

[1]Freeman Dyson, *The Origins of Life* (Cambridge: Cambridge University Press, 1985).
[2]Allen Newell, *Human Problem Solving* (Englewood Cliffs, N.J.: Prentice Hall, 1972).
[3]A. R. Luria, *Mind of the Mnemonist* (New York: Basic Books, 1968).
[4]Daniel W. Hillis, *The Connection Machine* (Cambridge: MIT Press, 1985).
[5]Stephen Wolfram, *Theory of Applications of Cellular Automata* (World Scientific, 1986).
[6]M. B. S. Haldane, *The Causes of Evolution* (Harper & Brothers, 1932).