Artificial Intelligence and Intelligent Design

Eric Scott, Andrews University

07 October, 2009

> *"The higher intellect, the imagination, the spirit, and even the heart might all find their congenial aliment in pursuits which, as some of their ardent votaries believed, would ascend from one step of powerful intelligence to another, until the philosopher should lay his hand on the secret of creative force and perhaps make new worlds for himself."* – Nathaniel Hawthorne, "The Birthmark"

As science has steadily given us more and more power over our world, the optimists have always hoped to unravel the secret behind the origin of life itself, in order that they might mend, imitate, and refine it, or at the very least gain a clue as to where we came from and what our place is in the cosmic metanarrative. Darwin is seen by most scientists as a bright beacon pointing to the answer. The Intelligent Design (ID) movement, popular among Christians in America though frowned upon by the scientific community, questions the sufficiency of natural processes to explain the profound complexity we see in nature, and furthermore declares that design by a higher intelligence can be inferred scientifically. Whether or not they denounce ID, scientists readily acknowledge that the precise mechanisms by which evolution has and does progress remain a mystery, and that Darwin, Watson, Crick, Gould etc. have provided only a vague outline of the story.

Those precise mechanisms, if they can be found, would be of extreme interest to technologists, and particularly in the software arena where programs can be written to match a well-defined model. Nature has solved some of the hardest problems we can imagine in the design of biological organisms. Artificial Intelligence (AI) can be briefly defined as the *science of solving hard problems automatically*, and thus it should be no surprise that recent advances in AI have come, not from the hard-coded and meticulously designed algorithms we traditionally associate with computer applications, but from biology-inspired "soft computing" frameworks that exploit our new and growing understanding of how complexity can self-organize in an autonomous way. The underlying concept is emergence, which has been revolutionizing virtually all the sciences over the last fifty years, as described by complexity theorists Ricard Solé and Brian Goodwin in their *Signs of Life: How Complexity Pervades Biology* (2000):

> *A remarkable burst of creativity in science is transforming traditional disciplines at an extraordinary rate, catalyzing movements whereby old boundaries are dissolving and newly integrated territories are being defined. The new vision comes from the world of complexity, chaos, and emergent order. This started in physics and mathematics but is now moving rapidly into the life sciences, where it is revealing new signatures of the creative process that underlie the evolution of organisms.*[1]

The same concepts that inspire soft computing, one of the most successful and promising branches of AI, fuel the optimist's hope that we are on the brink of discovering the secret behind life itself. By and large, ID has not confronted this growing paradigm. For the sake of broadening the discussion, therefore, in the below we use the perspective of emergence to evaluate two questions:

- A) Can we reliably infer design in a complex system?
- B) Does current theory provide a full explanation for the emergence of biological complexity?

Evolution as Function Optimization

Natural selection works by exploring a "fitness landscape" through mutations and sexual reproduction, looking for peaks which represent survival optimums in the current environment. The process works best when there is a smooth gradient leading upwards to a maximum, providing small steps that are acheivable with probabilistically feasible mutations (see Figure 1). Each correct mutation takes part of the population a step up the hill, providing the survival benefit and selective pressure necessary to carry the species to

[1] Ricard Soleé and Brian Goodwin, *Signs of Life: How Complexity Pervades Biology* (US: Basic Books, 2000), ix.

the top. All nature's algorithm need do is stumble into the so-called "basin of attraction" of one of these peaks. In this sense the evolutionary process can be seen as a stochastic (random) optimization algorithm, although we acknowledge that the dynamics of biological evolution in real-world ecosystems are more subtle than this simple caricature would imply. We stress the optimization interpretation because it can help in visualizing the limits of Darwinian processes.

Figure 1: *A fitness peak with a wide basin of attraction conducive to discovery via natural selection.*

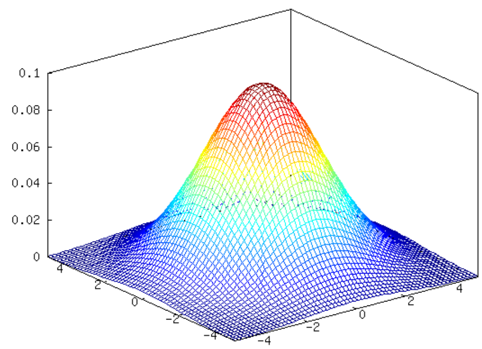

Constructing good solutions to real-world problems is more difficult than the smooth curve in Figure 1 makes it appear, which is why gradualistic models of evolutionary history have long been abandoned in favor of such ideas a punctuated equilibrium, in which progress occurs in occasional spurts. In practice all optimization algorithms suffer from the "local maximum" problem, as the fitness landscape is often tumultuous. The program happily climbs to the top of an easy-to-find, shallow peak, while altogether missing a much better solution with a narrow basin of attraction. There is no foreseen benefit to going towards the higher peak, since things don't get gradually better as we build that system (See Figure 2), and furthermore the valley between them may be too low for the organism to survive at all in an intermediate state.

Since a real fitness landscape has many more than two dimensions, going a long ways by chance and hitting the jackpot is extremely unlikely, and so a roadblock occurs in the evolutionary process. Intelligent Design focuses on instances where the local maximum problem appears to have been inexplicably transcended.

Figure 2: *A simulated fitness function with two local maximums, the taller of which has a very small basin of attraction.*

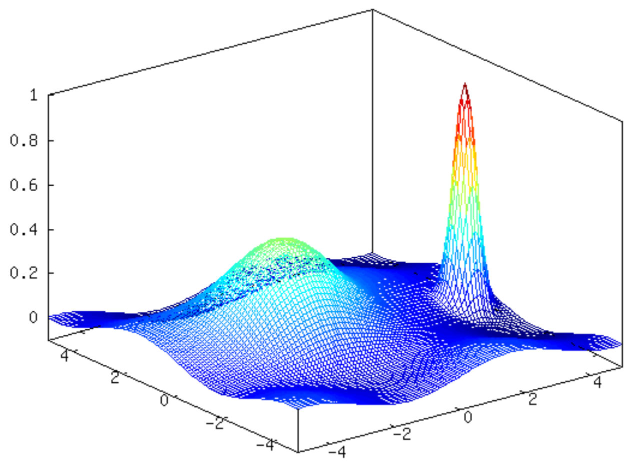

Design Inference

The premises that underlay ID are perhaps most clearly laid out by philosopher and mathematician William Dembski's three criteria for "design inference." The idea is that systems have been observed in nature which display the following properties, what Dembski calls Complex Specified Information (CSI):[2]

- **Contingency** requires that the system could have been built differently. It's like choosing Scrabble letters at random from a bag. If the bag contains only one type of letter, then our result is not contingent on anything, and we have no grounds to be surprised. However, if the bag contains a large variety of possibilities we may be able to infer design if a given selection can be shown to be unexplainable probabilistically.

[2]William Dembski, *Intelligent Design: The Bridge Between Science and Theology* (IL: Intervarsity Press, 1999), 128.

- **Complexity** by Dembski's definition requires that the system have many interdependent parts, without any of which the system would be incomplete.

- **Specification** requires that said system be well-suited for the problem we are attempting to solve. Out of the vast array of contingent possibilities, there are many "wrong" answers. In the Scrabble letters analogy, specification could be equivalent to drawing letters at random and trying to form a word in some particular language.

In applying these principles to biology, a system may provide survival value (specification), but with sufficient complexity it can be seen as impossible to have been produced via natural selection. Dembski sets the threshold information content of such an "irreducibly complex" system (a term coined by Michael Behe) at 500 bits, a number well large enough that it cannot be attributed to chance mutation.

If evolutionary processes are insufficient, it is reasoned, materialist explanations have failed us, and we can infer the presence of design. And, indeed, if a system unexplainable by incremental natural selection and requiring 500 bits of information to develop simultaneously and by chance can be shown to exist, it would be a compelling argument.

Hypothesis Testing

Analyzing the truth of a hypothesis such as Intelligent Design demands that we have a sound picture of the models and evidences involved, and how they affect our conclusions. That is, we need a clear basis for probabilistic inference – a framework which is absent from much of the ID literature.

In its quest to automate reasoning via precise computer programs, Artificial Intelligence maintains an endemic relationship with the endeavors of probability theorists who seek to quantify the process of reasoning from incomplete data. Most of AI applications are essentially mechanisms for generating and testing hypotheses and fitting new models to data. It is to our benefit, then, to adopt this paradigm of exactness in order to avoid the foggy misconceptions and fallacies that follow when we trust to intuition and the imprecise language which accompanies it. In a sense, we seek a *meta-model*, or a model of science itself, to assist communication and bring about a more perspicuous view of the issue at large.

The scientific method determines the relative plausibility of one model over another by testing their respective predictions. We can never prove that the model we have is *the* best one, the one and only truth, though we may become extremely confident in it after many tests (ex. Newton's law of gravity). The best we can do is show that its predictions match old and new data better than any other model we've been able to think of. Thus the proper epistemological context for the problem of origins is more subtle than a simple "science proves this" or "science proves that." Perhaps at this point it would be prudent to keep in mind what has been called the "eleventh commandment for statisticians:"

> *"Thou Shalt not Extrapolate Beyond the Range of Thy Data."*[3]

We will not cede to agnosticism, however, until the issue has been sufficiently explored.

Our strength of belief in a hypothesis is updated by evidence. As supporting evidence is accumulated, our confidence increases, or vice versa as detracting evidence is uncovered. The snafu is that our conclusion is partially determined by how confident we were to begin with – i.e. the initial "prior probability" we choose based on previous experience. This introduces an inherent subjectivity into our model of reasoning, as noted by eminent probability theorist E.T. Jaynes:

> *"So to our robot there is no such thing as an 'absolute' probability; all probabilities are necessarily conditional on [the prior probabilities]"*[4]

Some readers will recognize this approach – updating priors with evidence – as the *Bayesian* interpretation of statistical inference, which forms the foundation of a sizeable portion of traditional Artificial Intelligence tools. [5]

Many religions – occidental and oriental – emphasize personal experience as the primary source of evidence for their faith. Such evidence can be a valid part of the *personal* induction process – vastly affecting a person's priors, and thus conclusions – but is not amenable to the *public* analysis of

[3]Stephen P. Ellner and John Guckenheimer, *Dynamic Models in Biology* (Princeton: Princeton University Press, 2006), 2.

[4]E.T. Jaynes, *Probability Theory: The Logic of Science* (2003), 87-8.

[5]Note that I am not supporting relativism. One's priors are *always* based on previous evidence. Exactly how that evidence is converted to quantitative figures is where the subjectivity lies. Without evidence, we begin agnostic.

scientific discourse. Design Inference holds our interest specifically because it promises a public, scientific way to infer the existence of a deity.

Evaluating ID

What we have are two competing models for how life came to be which make the following broad statements and predictions, each of which are themselves hypotheses to be tested:

1. **Supernatural**:
   - A) Complexity exists that is difficult to explain with natural causes.
   - B) Complexity exists that has no natural explanation.
   - C) Some sort of supernatural entity designed this complexity.
2. **Natural**:
   - A) Simple *in vitro* and/or *in silica*[6] processes *can* produce a profound amount of superficial complexity.
   - B) Similar *in vivo* natural processes *do* produce some of the complexity we see in nature.
   - C) Similar natural processes produce *all* the complexity we see in nature.

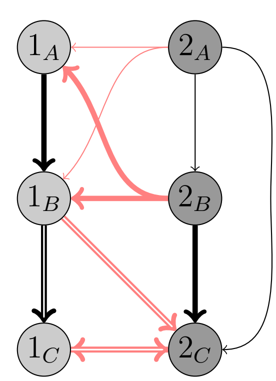

Figure 3: *Bayesian relationships between the hypotheses.*

Hypothesis 1(B), our focus, is a negative, and thus very difficult to prove. Its verification depends on an assurance as to the limits of *all* natural processes and *our ability to identify violations of those limits*, a prediction that can be supported by instances of evidence 1(A). 1(B) is not strictly falsifiable, but our confidence in it is reduced as we succeed in explaining more and more complex systems naturally, i.e. if 2(A) and (especially) 2(B) are verified in many cases. 1(C), which follows from the verification of 1(B), is only falsifiable by 2(C), which is equivalent to another negative (*no* complexity was created non-naturally) and thus all

[6] *In silica* refers to an experiment conducted solely via computer simulation.

but impossible to prove conclusively. 2(C) is falsifiable, however, by a *single* counterexample via 1(B). Figure 3 summarizes these observations in a rough version of what is known as a Bayesian network. Single lines represent supporting (black) or detracting (red) evidence, thicker lines having more effect than thin ones. Double lines represent logical proof/disproof.

We have two general sources of evidence in this model: apparent irreducible complexity, and its potential explanation as computational and physical experiments threaten to unveil simple explanations for an arbitrary level of apparent complexity.

Emergent Complexity

In seeking examples of irreducible complexity, our efforts are confounded by the difficulty in inferring the existence of "genuine" complexity to begin with. Yaneer Bar-Yam, director of the New England Complex Systems Institute, underscores this awareness in his text on complexity

> *"Complex phenomena require, by their nature, a complex model to generate them. This means we cannot expect simple models to generate truly complex behavior. Thus, a basic skepticism about the ability of theory to describe biological phenomena can be justified. What is missing, however, is an ability to know, a priori, what are truly complex phenomena and what properties of complex organisms can be attributed to simple universal behaviors."*[7]

Another complex systems researcher defends the value and relevance of *in silica* theory:

> *"One possible approach is through realistic models, that include as much detail as possible. On the other extreme, simple models with a minimum number of parameters allow for the determination of the basic ingredients necessary for the emergence of complex structures."*[8]

While examining simulations of complexity in various fields may not always tell us something for certain about natural systems, it does affect our picture

[7]Yaneer Bar-Yam, *Dynamics of Complex Systems* (CO: Westview Press, 1997), 690.

[8]Gomez Portillo IJ, Gleiser PM, "An Adaptive Complex Network Model for Brain Functional Networks", *PLoS ONE*, vol. 4 no. 9, 07 September, 2009.

of how apparent complexity may or may not arise from simple origins in a variety of situations.

Dembski's definition of complexity implicitly deals with the minimum information required to specify the system, not the high-level intricacy we actually observe. This is known as the Kolmogorov or "algorithmic" complexity of the system, and is an extension Shannon's information theory.[9] In this model certain systems can be compressed, described by an algorithm that can be expressed in a relatively small amount of data. For example, to transmit the string "ABABABABABABABAB" we needn't transmit each character individually, but can instead send "8*(AB)", meaning "AB" is repeated eight times. Or if we want to convey the base-1 message "1111111111111111111" it can be compressed to "10011" in base-2. Kolmogorov complexity is defined to be the shortest possible signal + algorithm pair that can be used to reconstruct the entire signal within the given computational environment (be it chemical or digital).

Among the most astounding results the study of complexity has brought us is the realization of the intricate patterns that can arise from simple specifications (i.e. low Kolmogorov complexity). Brief algorithms can be defined which, by repeating the same operation, self-organize or "emerge" into astounding structures which the viewer would assume have much more complex specifications. A classic example is the fractal pattern evident in the Mandelbrot set.

The idea of a "fractal" is somewhat ill-defined, but at its heart lies the concept of *self-similarity* – a property shared by many nonlinear and chaotic systems – in which a subsection of the system contains a scaled-down snapshot of the whole. This is readily apparent in the Mandelbrot set, in which the scale models are accompanied by beautiful swirls that look like the brush strokes of a talented artist. And yet no human being specified this painting: the entire system is described by the straightforward complex (as in imaginary numbers) polynomial:

$$z_{n+1} = z_n^2 + c$$

which is arguably the simplest conceivable nonlinear map in the complex plane.

At this point Dembski makes an important observation: it takes more than this simple equation to produce the famous picture.

[9] Bar-Yam, *Complex Systems*, 706.

Figure 4: *Images from deep within the Mandelbrot set and the corresponding Julia set.*

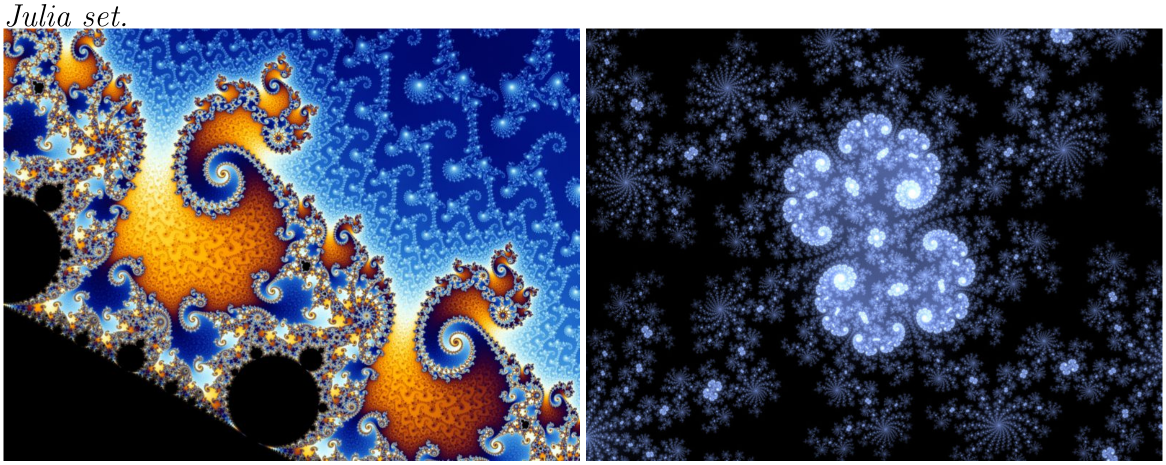

> *"Any function that produces a graphic depiction of the Mandelbrot set will be a complicated algorithm employing a complicated set of input data... But by itself the function $h(z) = z^2 + c$ is too information-poor to produce this graphic depiction of the Mandelbrot set j. Once we examine the precise informational antecedents to j, the illusion that we have generated information for nothing disappears."*[10]

In our initial examples of Komolgorov complexity, we didn't discuss how much information is required to specify the algorithms that decompress "8*(AB)" or convert base-2 back into base-1.

Clearly, however, as the length of the signal we want to reproduce increases, the algorithm size remains fixed. If the algorithm to convert base-2 into base-1, for example, takes 50 bits, then translating the number 10,000 in base-1 (10,000 bits) into base-2 ("1001110001000" – 13 bits) takes only 63 bits total. No matter the size of the algorithm, as the length of the signal increases its Kolmogorov complexity is significantly smaller than the original (except in the case of truly random data, which cannot be compressed). The Mandelbrot set is *infinitely* intricate, and its requisite algorithm is by no means large – Dembski describes the details in a single paragraph.

It is conceivable that algorithms of this sort tend to be irreducibly complex. It would indeed be difficult to encode the entire mechanism for the

[10]Dembski, 165

Mandelbrot set into under 500 bits (about 63 ASCII characters). It's important to note, however, that once any system is reduced to its fundamentals, our "algorithm" consists of the laws of nature themselves. The semantics that define the results of biological complexity operate according to these predefined physical principles; only the initial (compressed) input string need be provided.

A potent example of emergence in physics is found in the so-called "nanoflowers" that have been synthesized in laboratories under a variety of conditions. The structures apparent in Figure 5 were created by heating gallium nitride on a silicon substrate to 1100°C and then exposing the system to methane gas. "Interest in such structures," writes the Cambridge research team who developed the experiment, "centres around the combination of a simple growth process based on SiC nanowire formation, with a resultant structure having potentially complex mechanical and optical properties."[11]

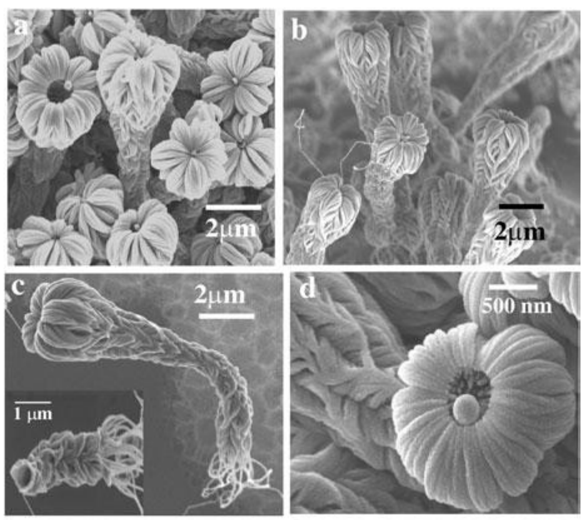

Figure 5: *These silicon carbide "nanoflowers" were inorganically synthesized via a simple chemical reaction.*

Neither the Mandelbrot set nor nanoflowers, impressive though they are, meet Dembski's *specification* criterion, and thus are not sufficient for demonstrating the significance of emergence in biology. The important idea here,

[11] Ghim Wei Ho et al., "Three-dimensional crystalline SiC nanowire flowers," *Nanotechnology*, 15 (2004), 996-999.

however, is that *patterns which seem complex to our eye can actually emerge from a simple process* – that is, they have a low Kolmogorov complexity – and that is significant to our purposes.

Stephen Wolfram, whose developments in computer science have made significant contributions to our understanding of complexity, points out that

> *"We have seen a great many systems whose underlying rules are extremely simple, yet whose overall behavior is sufficiently complex that even by thinking quite hard we cannot recognize its simple origins."*

Wolfram goes so far as to consider the case of trying to infer design in an extraterrestrial radio signal (also an illustration Dembski uses to discuss ID[12]), intimating that awareness of emergence causes *any* seemingly designed piece of information fall suspect to natural causes. "We cannot find an abstract way to give evidence of purpose or intelligence," he says.[13]

Whether or not we take such a laconic position as Wolfram, it should be clear that *design inference is made difficult by the tendency for complex systems to display emergent properties.* Furthermore, this self-organization makes it *easier* for natural selection to explain biological systems. If a system can be specified with a small number of bits, then the local maximum problem becomes a non-issue. Something as superficially intricate as the Mandelbrot set can be defined by a very simple mathematical algorithm, and so it is conceivable that something such as an eye could be described by a simple gene network.

In this view, natural selection chooses apparently complex specified systems only from amongst emergent designs that are feasible to discover via chance and gradual development; each new available step builds on the matrix of interacting parts available given the organism's current configuration (This process by which an organism's current state affects it's future possibilities has been called *autocatalysis*[14]), causing gradual evolutionary steps to quickly produce extremely intricate complexity. Indeed emergence must take place at least some extent, since our genome is many orders of magnitude too short to define every feature of our physiology explicitly.

[12]Dembski, 128

[13]Stephen Wolfram, *A New Kind of Science* (IL: Wolfram Media, 2002), 620, 838.

[14]Russell Eberhart and Yuhui Shi, *Computational Intelligence: Concepts to Implementations* (Amsterdam: Elsevier, 2007), 28.

We have only scratched the surface of the nonlinear mathematics, computing, and physics that explores *in silica* and *in vitro* models of emergence (our hypothesis 2(A)). Clearly, however, such things are only the first step towards understanding and uncovering emergence in real biological systems (2(B)). In the last decade research seeking empirical confirmation of the predictions of emergent models in the real world has placed a great deal of emphasis on power-law distributions which are easy to detect statistically. Complexity researcher Melanie Mitchell summarizes the source of interest:

> *"Power-law distributions have been identified for the size of cities, people's incomes, earthquakes, variability in heart rate, forest fires, and stock-market volatility, to name just a few phenomena... There are many different explanations of power laws observed in nature (e.g. preferential attachment, fractal structure, self-organized criticality, highly optimized tolerance, among others), and little agreement on which observed power laws are caused by which mechanisms."*[15]

What the true mechanisms are, as well as how to detect them, remain largely a mystery.

Emergent Computing

One complexity researcher cites technological applications as the primary motivation for interest in the field:

> *"At the core of this explosion of interest is the realization that both natural and artificial systems (mostly computer models) are both quite capable of showing several complex phenomena in common."*[16]

"Soft computing" is where the paradigm of emergence meets technology. The term, coined by Lotfi Zadeh in 1994[17], refers generally not only to probabilistic AI tools, but especially to nature-inspired systems such as artificial neural networks, evolutionary computation, swarm intelligence, and artificial

[15] Melanie Mitchell, *Complexity: A Guided Tour* (Oxford University Press, 2009), 269.

[16] Octavio Miramontes, *Complexity and Behavior in Leptothorax Ants* (2007).

[17] Lotfi A. Zadeh, "Fuzzy Logic, Neural Networks, and Soft Computing," *Communications of the ACM*, March 1994, Vol. 37 No. 3, 77-84.

immune systems[18]. Each of these technologies is fuelled by the expectation that self-organizing, emergent properties can bring about sophisticated solutions to our problems even when humans are at a loss at how to approach them. The paradigm has struck the field deeply enough that some advocate teaching the theory of computation as a property of the natural world.[19]

In general, soft computing algorithms are machine learning algorithms. That is, they are applied to a library of sample data, modify and/or parameterize themselves accordingly, and then are tested on other data to see how well they've modeled the system in question. The problem to be solved might be object recognition, optimization of an industrial process, spam detection, or modeling the dynamics of a nonlinear mechanical system being produced for use in vehicles. All these methods resemble biological evolution in that it is not just the algorithm and initial state of the system that produce the complexity, but the information that is generalized and teased out of the environment via selective processes.

We can think of these applications as a sort of laboratory to test our understanding of the natural mechanisms that might give rise to specified complexity. If we genuinely understand how biological solutions arise, then we should be able to emulate those processes to solve our problems. Indeed these attempts have allowed us to solve profoundly difficult problems and develop practical solutions that are otherwise intractable. The evolved solutions are furthermore often too complex to reverse engineer, and remain a black box to their creators. Melanie Mitchell summarizes the difficulty:

> *"The key [to artificial evolution], it turns out, is not [isolated] genes, but the way different genes interact, just as has been found in real genetics. And just as in real genetics, it's very difficult to figure out how these various interactions lead to the overall behavior or fitness."*[20]

Despite the success of these techniques in generating complexity, artificial neural networks are far, far off from the power of a biological brain, and artificial immune systems do not display the sort of powerful adaptability evident in their human counterpart. The evolutionary approach *in silica* is

[18] Andries P. Engelbrecht, *Computational Intelligence: An Introduction (England: John Wiley & Sons, Ltd, 2007)*, 3.

[19] Colin G. Johnson, "Teaching Natural Computation," *IEEE Computational Intelligence Magazine*, Feb. 2009, Vol. 4 No. 1, 24-30.

[20] Mitchell, 136

not a magic formula that can relieve the engineer from his responsibilities as designer, as Bar-Yam points out:

> *"While the GA/EA [genetic/evolutionary algorithms] approach can help in specific cases, it is well known that evolution from scratch is slow. Thus it is helpful to take advantage of the capability of human beings to contribute to the design of systems."*[21]

The general models we have for how complexity emerges in specified ways have not provided us, so far, with the detailed mechanistic understanding required to implement comparable artificial solutions. This signals us that we have more to learn, and that our models still require a lot of fleshing out before we can pretend that we understand the nature of high-level reality.

Conclusion

In this essay we set out to answer the questions A) Can we reliably infer design in a complex system? And B) Does current theory provide a full explanation for the emergence of biological complexity? The answer to both, we conclude, is *no*.

Design inference, we said, requires an assurance as to our ability to identify violations of the limits of natural selection, places where the local maximum problem has somehow been significantly transcended by a 500-bit leap in the dark. Dawkins laconically calls the identification of irreducibly complex systems the "argument from personal incredulity"[22]. In our exploration of *in silica* and *in vitro* complexity, we see that the limits of natural selection are indeed by no means easy to recognize. We simply have no method of confidently establishing the inherent complexity of a system. As such, whatever evidence there is for design is at present insubstantial, as our "incredulity" could well be unfounded.

It is the author's conclusion, then, that Intelligent Design does not stand on its own two feet as a public scientific theory. Nothing in this essay, however, has proved that a deity was never involved in the history of life. Judea Pearl, the father of probabilistic AI, has formally pointed out the distinction

[21] Yaneer Bar-Yam, "When Systems Engineering Fails – Toward Complex Systems Engineering," *Conference on Systems, Man, and Cybernetics* 2003 Vol. 2 (IEEE Press, Piscataway, NJ, 2003) pp. 2021-2028. Available online at http://www.necsi.edu/research/engineering/

[22] Richard Dawkins, *The God Delusion* (2006), 128.

between *sufficient* and *actual* causes[23]: simply because natural processes might be sufficient for producing life does not prove that they were the actual cause of all of it. Intelligent Design is plausible, then, only when an individual's prior experience leads them to be inclined towards belief in a deity who is actively involved in the present and past development of life. That is, it can be a valid part of a religious world view already established via personal experience or some other evidence, but *ID cannot be used to prove God*, per se.

As to the second question, have we provided a full *natural* explanation for biological complexity? The small amount of empirical evidence supporting various theories of emergence only gets us so far in affirming our models. Furthermore, the limited capabilities of our technological imitations of nature demonstrate how much we still have to learn. Human design is still integral to the process, and it may still remain so for a long time even if complexity does hold the answer.

If one's prior experience leads him or her to fairly weak evidence in one way or another, when should he or she adopt an opinion and/or remain agnostic? We'll give Jaynes the final word on the matter:

> *"There is nothing in probability theory* per se *which can tell us where to put these critical levels at which we make our decisions. This has to be based on value judgments: what are the consequences of making wrong decisions, and what are the costs of making further tests?"*[24]

Recommended Reading

There is a host of accessible literature written about complexity and its sister fields by scientists of great renown. Below are works that may be useful to those interested in learning more about the movement at an introductory or avocational level.

- Albert-László Barabási, *Linked: How Everything Is Connected to Everything Else and What It Means for Business, Science, and Everyday Life* (2003). Social and scientific complexity as modelled by small-world and scale-free networks.

[23] Judea Pearl, *Causality: Models, Reasoning, and Inference* (New York: Cambridge University Press, 2000), 309.

[24] Jaynes, 96.

- Nancy Forbes, *Imitation of Life: How Biology is Inspiring Computing* (2004). The story of soft computing.

- James Gleick, *Chaos: Making a New Science* (1987). The best-selling exposition of chaos theory's history and its relation to complexity.

- Douglas Hofstader, *Gödel, Escher, Bach: An Eternal Golden Braid* (1979). The Pulitzer-winning masterpiece on the emergent nature of intelligence that has inspired a generation of scientists and mathematicians.

- John Holland, *Hidden Order: How Adaptation Builds Complexity* (1996). Complexity as explained by the father of genetic algorithms. Considerably technical.

- Stuart Kauffman, *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity* (1994). Inspiringly optimistic (if speculative) catalog of how complexity can solve our hard problems and show that life is not as improbable as one may suppose.

- John H. Miller and Scott E. Page, *Complex Adaptive Systems: An Introduction to Computational Models of Social Life* (2007). Agent-based computer models applied to social problems. Discusses the philosophy of modelling and how much faith we should invest in computer simulations.

- Melanie Mitchell, *Complexity: A Guided Tour* (2009). A captivating, accessible, and thorough summary of complexity theory and related fields. One of the best general introductions I've found for casual reading.

- Richard Solé and Brian Goodwin, *Signs of Life: How Complexity Pervades Biology* (2000). A beautiful exposition of the various aspects of complexity theory and how they relate to biology. Mathematics included, but optional and boxed off from the main prose.

- Steven Strogatz, *Sync: How Order Emerges from Chaos in the Universe, Nature, and Daily Life* (2003). Complexity as approached from the theory of coupled oscillators and nonlinear dynamics.

- Stephen Wolfram, *A New Kind of Science* (2002). Complexity as unveiled by cellular automata.