In Dmitry Bulatov, ed., *Biomediale: Contemporary Science and Genomic Culture* (pp. 216-231). Kaliningrad: The National Center for Contemporary Arts, 2004.

Artificial life illuminates human hyper-creativity

Mark A. Bedau
Department of Philosophy, Reed College, Portland OR 97202, USA
mab@reed.edu
http://www.reed.edu/~mab

The aim of this chapter is to show how the technological research activity called "artificial life" is shedding new light on human creativity. Artificial life aims to understanding the fundamental behavior of life-like systems by synthesizing that behavior in artificial systems (more on artificial life below). One of the most interesting behaviors of living systems is their creativity. Biological creativity can be found in both individual living organisms and in the whole biosphere—the entire interconnected system comprised of all forms of life—but I will focus in this chapter on the biological creativity exhibited by the evolutionary process. This is the creativity that enabled the earliest simple life forms to spontaneously evolve into the incredibly rich and beautiful diversity of life that now surrounds us. This diversity of life includes the most complex adaptive and intelligent systems in the known universe. This is an amazingly powerful spontaneous creation process, indeed. I will refer to it as *hyper-creativity* to call attention to the way in which it produces *qualitatively new and more complex* kinds of adaptations. There is a similar quality in human creativity. I am thinking of the aesthetic and cultural creativity of artists, but also the intellectual creativity of scientists and scholars, as well as the commercial and practical creativity of craftsmen, businessmen, and entrepreneurs. And I want to focus especially on the *hyper-creative* aspects of human creativity—the way in which human activity can yield *qualitatively new and more complex* creations.

Artificial life illuminates human creativity by means of a simple two-stage argument from analogy. The first stage is that artificial life illuminates biological creativity, and the second stage is that there is a deep analogy between human and biological creativity. These two stages yield the conclusion that artificial life's insights into biological creativity analogously illuminate human creativity.

Artists are quite inventive and perceptive about exploiting the potentials of new technologies for their aesthetic purposes. Artificial life technology is no exception. There are at least three ways in which contemporary artists are using artificial life. First, technologically savvy artists can deploy artificial life technology for a variety of aesthetic purposes. Examples start with using artificial life methods and insights to produce new and better computer animations of life forms, and examples extend to those who use artificial life's bottom-up evolutionary and generative processes to create new kinds of active art, evolving art, and interactive art; examples include Bilotta et al. (2000), Sommerer and Mignonneau (2000), Mignonneau and Sommerer (2001), and

Innocent (2001). Second, artificial life is radically changing human culture and technology, and this trend will expand in time. Art often responds to and comments on such changes, and this is already happening with the changes wrought by artificial life; for example, see Huws (2000) and Ando (2001). Third, art has a long tradition of representing and responding to our understanding of nature. Artificial life is adding a distinctive perspective to our view of nature, and these insights are sparking new aesthetic objects; three recent instances are Youngs (2000), Prophet (2001), and Whitelaw (2001).

The traffic between artificial life and art travels in both directions. Not only do artists use artificial life for their own purposes; they also promote scientific progress in artificial life in various ways. For one thing, artists that use artificial life techniques and insights can be viewed as the "consumers" of the "product" that artificial life produces, and consumer demand is always a spur to producing better products. Scientists also gain a broader perspective on their own scientific activity when artists explore the implications of the science and subject it to commentary and social criticism in their distinctive way. Finally, human aesthetic activity is one of the most striking and distinctive manifestations of the creative potential of living systems. So, it would behoove those who want to understand life's creative potential to keep an eye on the latest developments within the arts. This last issue is closely related to the topic of the present chapter: how artificial life illuminates the human creative process, including that exhibited by artists. The issue here is our *understanding* of the human creative process.

Overview of artificial life

Let me first explain what artificial life is. Artificial life can be situated within an interdisciplinary activity devoted to understanding the behavior of complex systems. Examples of this new venture include the science of chaos (Crutchfield et at. 1986) and studies of cellular automata (Wolfram 1994; Langton 1992; Wolfram 2002). By abstracting away from the details of chaotic systems, one can discern fundamental properties that unify and explain a diverse range of chaotic systems. Similarly, by abstracting away from the details of life-like systems and synthesizing these processed in artificial media, typically computers, the field of artificial life seeks to understand the essential processes shared by broad classes of life-like systems. While biology aims to understand life-as-we-know-it, artificial life's interest extends to all of life-as-it-could-be (Langton 1989).

It is useful to contrast artificial life with its well-known sister field: artificial intelligence (AI). Roughly speaking, what AI is to psychology, artificial life is to biology. Both focus on computational systems but AI is interested in systems that produce cognitive processes such as reasoning, memory, and perception while artificial life is interested in systems that produce the processes characteristic of living systems. These processes include the spontaneous generation of order and self-organization, self-reproduction and autonomous adaptation, and open-ended evolution.

Despite these similarities, there is an important difference between the modeling strategies artificial intelligence and artificial life typically employ. Most traditional AI models are top-down-specified serial systems involving a complicated, centralized controller that makes decisions based on access to all aspects of global state. The controller's decisions have the potential to affect directly any aspect of the whole system. On the other hand, most natural systems exhibiting complex autonomous behavior seem to be parallel, distributed networks of low-level communicating “agents.” Each agent's decisions is based on information about only the agent's own local state, and its decisions directly affect only its own local situation. Following this lead, artificial life is exploring forms of emergent global order produced by bottom-up-specified parallel systems of simple local agents. Not only do artificial life models share the bottom-up architecture found in natural systems that exhibit complex autonomous behavior, but the flexible “intelligent” behavior that spontaneously emerges from artificial life models is also strikingly akin to that found in nature. Thus, artificial life models share some important features with the distributed, bottom-up connectionist models that have recently revolutionized AI (Rumelhart and McClelland 1986). (For a discussion of some important differences between artificial life and connectionism, see Bedau 2002a.)

The bottom-up architecture of artificial life systems allows micro-level entities to continually affect the context of their own behavior. This allows artificial life systems to capture some of the spontaneous creativity inherent in living systems. For example, a population of organisms typically has an active hand in constructing the environment to which it adapts (Bedau, 1996). Because of the network of interactions among organisms, an organism's adaptation to its environment typically changes the intrinsic properties of the external objects in its environment. Nevertheless, in order to insure mathematical tractability, all too many models of organisms within an environment ignore these interactions. Such interactions would imply a population of entities “undergoing a kaleidoscopic array of simultaneous nonlinear interactions”, as John Holland puts it (Holland 1992, p. 184). The only way to study the effects of these interactions is to do what the field of artificial life does: build bottom-up models and then empirically investigate their emergent global behavior through computer simulations.

Artificial life models routinely do show impressive global phenomena emerging from simple micro-level interactions. Flocking behavior is one vivid example of this. Flocks of birds exhibit impressive macro-level behavior. The flock maintains its cohesion while moving ahead, changing direction, and negotiating obstacles. And these global patterns are achieved without any global control. No individual bird issues flight instructions to the rest of the flock; no central authority is even aware of the global state of the flock. The global behavior is simply the aggregate effect of the microcontingencies of individual bird trajectories.

Natural flocking behavior can be feasibly produced by Craig Reynolds’s “Boids” system (1987, 1992; Reynolds created Java demo of the Boids that is available on the web at http://www.red3d.com/cwr/boids/). When one views Reynold's flocking demos, one is vividly struck by how natural the flocking behavior seems. The collection of individual boids spontaneously organize into a flock that then maintains its cohesion as it moves and changes direction and negotiates obstacles, fluidly flowing through space and time. The flock is a loosely formed group, so loose that individual boids sometimes lose contact with the rest of the flock and fly off on their own, only to rejoin the flock if they come close enough to the flock's sphere of influence. The flock appropriately adjusts its spatial configuration and motion in response to internal and external circumstances. For example, the flock maintains its cohesion as it follows along a wall; also, the flock splits into two subflocks if it runs into a column, and then the two subflocks will merge back into one when they have flown past the column.

The Boids system produces these natural, supple flocking dynamics as the emergent aggregate effect of micro-level boid activity. Each boid acts independently in the sense that its behavior is determined solely by following the imperatives of its own internal rules. An individual boid's dynamical behavior affects and is affected by only certain local features of its environment—nearby boids and other nearby objects such as walls and columns. The Boids system contains no explicit directions for flock dynamics. The flocking behavior produced by the system consists of the aggregated individual boid trajectories and the flock's global dynamics emerges out of the individual boid's explicit micro-level dynamics.

Reynold's Boids provides one illustration of how complex phenomena of living systems can emerge from simple bottom-up artificial life systems. This pattern has many other instances. Consider one more example: evolution, which is one of the hallmarks of living systems (Bedau 1996a). One might speculate indefinitely about the minimal conditions and ultimate potential of such a process, but a feasible model can cut through such speculation. Tom Ray's (1992) Tierra is such a model. Tierra consists of a population of self-replicating machine language programs that "reside" in computer memory consuming the "resource" CPU time. A Tierran "genotype" consists of a specific type of string of self-replicating machine code, and each Tierran "creature" is a token of a Tierran genotype. A simulation starts when the memory is inoculated with a single self-replicating program, the "ancestor", and then left to run on its own. At first the ancestor and its off-spring repeatedly replicate until the available memory space is teeming with creatures which all share the same ancestral genotype. However, since any given machine language creature eventually dies, and since errors (mutations) sometimes occur when a creature replicates, the population of Tierra creatures evolves. Over time the "ecology" of Tierran genotypes becomes remarkably diverse, with the appearance of fitter and fitter genotypes, parasites, and hyper-parasites, among other things.

By exploring the behavior generated by specific bottom-up systems like Reynold's Boids and Ray's Tierra, the field of artificial life studies how the global phenomena characteristic of living systems can spontaneously emerge from the interactions among simple micro-level agents. By illuminating the minimal conditions sufficient to produce these phenomena, the artificial life systems help us to understand not only how such phenomena happen in the actual world but also how they could happen in any possible world.

Artificial life and biological creativity

Artificial life has a unique ability to shed light on the nature of life's "hyper creativity". This illumination to date is partly negative; it shows what is wrong with the prominent contemporary perspectives on life's creativity. A good illustration of this negative argument comes from assessing the debate between Gould and Dennett on long-term evolutionary trends.

The progression of evolution in our biosphere exhibits a remarkable overall increase in complexity. From simple prokaryotic one-celled life evolved eukaryotic cellular life forms with a nucleus and numerous other cytoplasmic structures. From these evolution produced life forms composed out of a multiplicity of cells. Out of multicellular life evolved large-bodied vertebrate creatures with sophisticated sensory processing capacities. And from those beings evolution produced highly intelligent creatures that use language and develop sophisticated technology, i.e., humans. How should we think about this evolutionary trajectory that leads from simple life forms to those that are remarkably complex?

One possible explanation of the way life has evolved is the hypothesis that open-ended evolutionary processes have an inherent tendency to create creatures with increasingly complicated functional organization. Consider an analogy with thermodynamics. The second law of thermodynamics is an "arrow of entropy." It asserts that the entropy in physical systems has a general tendency to increase with time. Similarly, the biological hypothesis of the "arrow of complexity" asserts that the complex functional organization of the most complex evolved organisms has a general tendency to increase with time.

The fact that the evolution of life is consistent with the arrow of complexity hypothesis does not establish the truth of the hypothesis, of course. Stephen Jay Gould has been an especially vigilant guard against any the idea that evolution embodies any form of progressive trend. In his book *Wonderful Life* (1989) on the fossils in the Burgess shale, Gould explains how the evolution of life can be understood as a process free from anything like the arrow of complexity. The book's central argument is that anything that looks like an evolutionary progression is really just a contingent by-product of myriad accidents frozen into the evolutionary record.

Gould's argument appeals to his view of the essential contingency of historical processes like evolution. Gould thinks that the contingency of the evolution is inconsistent with general laws like the arrow of complexity. The results of historical processes "do not arise as deducible consequences from any law of nature; they are not even predictable from any general or abstract property of the larger system...." (p. 284). Instead, "almost every interesting event of life's history falls into the realm of contingency" (p. 290).

Gould illustrates his argument with a brilliant thought experiment that he calls “replaying the tape of life.” This thought experiment involves imagining that the history of the evolution of life were recorded on a tape. We imagine that we rewind the tape backward in time, erasing the evolutionary process, and then we play the tape forward again, but this time we allow different accidents, different historical contingencies, to leave their mark on the evolution of life.

> You press the rewind button and, making sure you thoroughly erase everything that actually happened, go back to any time and place in the past—say, to the seas of the Burgess Shale. Then let the tape run again and see if the repetition looks at all like the original. If each replay strongly resembles life's actual pathway, then we must conclude that what really happened pretty much had to occur. But suppose that the experimental versions all yield sensible results strikingly different from the actual history of life? What could we then say about the predictability of self-conscious intelligence? or of mammals? or of vertebrates? or of life on land? or simply of multicellular persistence for 600 million years? (pp. 48-50).

Gould thinks that this thought experiment refutes the arrow of complexity, for “any replay of the tape would lead evolution down a pathway radically different from the road actually taken” (p. 51). If Gould is right, then there is no inherent tendency for open-ended evolution to produce complexity and the arrow of complexity hypothesis is false.

But it is not clear that Gould *is* right. For example, Daniel Dennett (1995) draws exactly the opposite conclusion from the very same thought experiment. Dennett argues that certain complex features like sophisticated sensory processing provide a distinct adaptive advantage in a wide range of environments. Thus, natural selection will almost inevitably discover significantly advantageous features that are accessible from multiple evolutionary pathways. Examples of evolutionary convergence, such as flight and eyesight, illustrate this argument. Dennett concludes that replaying life’s tape will almost inevitably produce highly intelligent creatures that use language and develop sophisticated technology.

So, which conclusion *does* the thought experiment support? I think that replaying life's tape is the perfect experiment for testing the arrow of complexity hypothesis. But neither Gould nor Dennett shows any interest in pursuing the thought experiment in a concrete and constructive way. Their guesses about what would happen if you replayed the tape of life are just that—guesses, made *a priori* in the absence of empirical evidence about what the thought experiment would actually produce, and so reflecting little more than their respective antecedent biases about the situation.

Perhaps Gould and Dennett think the thought experiment is no more than a rhetorical exercise for projecting one’s favored perspective on the issues. But that would be a mistake. The thought experiment of replaying the tape of life is exactly the sort of investigation carried out time and again in artificial life. A standard methodology in artificial life is to create a system that embodies certain life-like properties of interest and then extensively study the system's behavior as system parameters and contingencies like random mutations are varied. You can think of these artificial life experiments as "road tests" designed to ferret out the system's typical behavior under all sorts of conditions. A thorough road test typically reveals a very detailed picture what happens when you "replay" the system.

One of the lessons learned from artificial life's extensive experience with such road tests is that our *a priori* expectations about their outcomes are highly fallible. The only sure way to determine what will happen is to create the relevant system and then observe its typical behavior. As I've said, artificial life is exactly where this sort of investigation occurs. A central goal of artificial life is to discover the inherent trends in evolving systems by devising systems that exhibit open-ended evolution (Bedau et al. 2000). With such a system in hand, you could rerun the tape of life to your heart's content. The detailed course of evolution in each instance would reflect the history of accidents unique to each it, but a general pattern behind all these contingencies could still emerge. Judicious analysis of the mass of contingencies collected from extensive road tests of the right systems would reveal whether an arrow of complexity lurks inside open-ended evolutionary processes in general. The best evidence in favor of the arrow of complexity hypothesis would come from showing that a tendency toward increasing adaptive complexity is the norm in such artificial life systems. Actually conducting these thought experiments introduces much needed discipline into the discussion. We can gain confidence that we understand how to explain some phenomenon only when we can synthesize a system that exhibits that phenomenon. When we are unable to do this, we simply reveal our ignorance. All conjectures about the arrow of complexity will remain up in the air until one creating and empirically observing the relevant thought experiments.

However, the fact of the matter is that *no one* has yet successfully replayed life's tape. It is not that people have not been trying. They have been trying, but they have been failing. The problem is that no one has been able to create a system that exhibits continual open-ended evolution of adaptive complexity. No one knows how to design a system that exhibits the kind of open-ended evolution characteristic of our biosphere. A number of artificial life systems have been advertised as plausible initial candidates for such systems, such as Tierra (Ray 1992), discussed above. But upon closer inspection, they have been proven to lack the kind of hyper-creativity observed in the biosphere. This negative result has been one of artificial life's most salient contributions so far to our understanding of biological creativity.

The method behind this result involves measuring the creative power of adaptive evolution. (See Bedau and Packard 1992, Bedau 1995, Bedau 1996, Bedau and Brown 1997, Bedau, Snyder, Brown, and Packard 1997, Bedau, Snyder, and Packard 1998, Bedau, Joshi, and Lillie 1999, Rechtsteiner and Bedau

1999a,b.) This method reveals that evolving systems fall into four qualitatively different classes. Class 1 consists of systems in which evolution creates no adaptations at all. Systems in which evolution has created adaptations but in which no new adaptations are being created fall into class 2. Class 3 consists of systems that continually create new adaptations but are bounded in the amount of adaptive structure they contain. If new adaptations are continually created and the total amount of adaptive structure continues to grow, then the system falls into class 4. The biosphere as reflected in the fossil record exhibits class 4 dynamics.

Class 4 is an especially explosive kind of evolutionary creativity. What I have been calling hyper-creativity would fall into class 4. What makes class 4 especially intriguing is that no known existing artificial evolving system generates class 4 behavior. All the artificial life systems that show the most promise of exhibiting open-ended evolution, such as Tierra, have been tested, but none shows the kind of creative evolution that we can see in the biosphere (Bedau, Snyder, Brown, and Packard 1997; Bedau, Snyder, and Packard 1998). These are the details behind artificial life's negative result concerning biological hyper-creativity.

However, there is a positive face of this result. Although we do not know the mechanism behind class 4 behavior, our results hint at what existing systems are missing. Specifically, hyper-creativity seems to arise when evolution continually creates new niches that open the door to qualitatively new kinds of adaptations. That is, the key seems to be the creation of special innovations that enable a family of qualitatively new kinds of adaptations to exist. One example is the innovation that led to the creation of multicellular life. Multicellular life has many new kinds of adaptive strategies open to it, such as the complex morphologies and cellular specializations that can then exist. Another example is the innovation that allowed life to colonize the land. Terrestrial life can have a host of new kinds of adaptations, such as those involving different forms of locomotion and flight. So, although no current artificial system exhibits hyper-creativity, we have distinct clues about where to look for such systems. And we also have a practical tool for gauging how well our new creations measure up to the example set by the biosphere's evolutionary creativity. In my opinion, it is just a matter of time before this particular scientific hurdle has been surmounted. And artificial life is the arena in which this contest is taking place.

The analogy with human creativity

I think that there is an analogy between human and biological creativity. In particular, I think that the culture of human artistic creations is analogous to an evolving population of biological organisms. The suggestion that human creativity shares a deep similarity with biological creativity has been voiced before, but artificial life can shed new constructive light on this analogy.

The current state of the art of the biological perspective on human creativity is exemplified by Daniel Dennett's recent Presidential Address to the American

Philosophical Association (Dennett 2001). Dennett's central thesis is that "all works of human genius can be understood in the end to be mechanistically generated products of a cascade of generate-and-test algorithms." What this means, in essence, is simple. This "Darwinian theory of creative intelligence," as Dennett calls it, construes the human creative process as exactly analogous to the creative process of evolution. Darwin's great insight into evolutionary creativity was that the process of natural selection spontaneously creates well-adapted organisms. That is, if you have a population of self-reproducing entities with varied features that are inherited from generation to generation, and if the nature of an entity's features can increase or decrease the chances that the entity will survive or reproduce, and if new varieties of features are somehow generated (perhaps randomly) with a suitable frequency, then over time the entities in the population will tend to acquire more and better adaptive features.

The analogous process for human creativity posited by the Darwinian theory of creative intelligence would be quite similar. Human culture creates and sustains a population of more or less abstract entities; they include ideas, hunches, insights, beliefs, methods, styles, procedures, and they were collectively lumped together in the category of "memes" by Richard Dawkins (Dawkins 19XX). Memes can enter our consciousness and influence our behavior, and they can be transmitted from person to person through conversation, popular media, or any other form of human communication. Memes have various features; some are memorable, some tend to cause altruistic behavior, some are understandable only by those with certain technical training, etc. Those features can affect the likelihood that a meme will be entertained, remembered, and communicated. Furthermore, new variants of existing memes are produced from time to time; the innovations can be either accidental (e.g., the result of misremembering) or purposeful (e.g., the result of conscious effort to modify a meme to perform some function better). These memetic innovations might increase or decrease the meme's chances of flourishing in the meme pool. So, putting all this together, over time the process of natural selection will tend to alter the memes in ways that make them more likely to flourish. Dennett's "generate-and-test algorithms" in the quote above refers to the process by which memetic innovations are generated and then the better-adapted memes are preferentially selected and preserved. He summarizes this Darwinian perspective on human creativity as follows:

> What process could conceivably yield such improbable "achievements of creative skill" [like Shakespeare's *Hamlet*]? What Darwin saw is that design is always both valuable and costly. It does not fall like manna from heaven, but must be accumulated the hard way, by time-consuming, energy-consuming processes of mindless search through "primeval chaos," automatically preserving happy accidents when they occur. This broadband process of Research and Development is breathtakingly inefficient, but—this is Darwin's great insight—if the costly fruits of R and D can be thriftily conserved, copied, and re-used, they can be accumulated over time to yield "the achievements of creative skill."

> "This principle of preservation," Darwin says, "I have called, for the sake of brevity, Natural Selection."

This Darwinian perspective on human creativity threatens many of our cherished views about our selves and our responsibility for what we create, and Dennett takes pains to respond to these doubts.

> There is a persistent problem of imagination management in the debates surrounding this issue: people on both sides have a tendency to underestimate the resources of Darwinism, imagining simplistic alternatives that do not exhaust the space of possibilities. … [A]nti-Darwinians, noting the huge distance between a beehive and the *St. Matthew Passion* as created objects, are apt to suppose that anybody who proposes to explain both creative processes with a single set of principles must be guilty of one reductionist fantasy or another: "Bach has a gene for writing baroque counterpoint just like bees' gene for forming wax hexagons" or "Bach was just a mindless trial-and-error mutator and selector of the musical memes that already flourished in his cultural environment." Both of these alternatives are nonsense, of course, but pointing out their flaws does nothing to support the idea that ("therefore") there must be irreducibly *non-Darwinian* principles at work in any account of Bach's creativity. In place of this dimly imagined chasm with "Darwinian phenomena" on one side and "non-Darwinian phenomena" on the other side, we need to learn to see the space between bee and Bach as populated with all manner of mixed cases, differing from their nearest neighbors in barely perceptible ways, replacing the chasm with a traversable gradient of non-minds, protominds, hemi-demi-semi minds, magpie minds, copycat minds, aping minds, clear-pastiche minds, "path-finding" minds, "ground-breaking" minds, and eventually, genius minds.

I find the picture painted by Dennett attractive. But it is unclear what kind of generate-and-test mechanism could produce the wonders of human creativity. Dennett's explanation is a compelling story, but the problem is that *it is just a story*. It is too vague to test constructively because Dennett does not specify a concrete generative system that we could actually study.

One might wonder whether this sort of constructive test is really necessary. Dennett suggests that the Darwinian theory does not need to be tested because it is "the only solution in sight." But that is just a confession of our current ignorance rather than a proof that his story is correct. Consider an analogy. Before Darwin came up with his explanation, someone like Paley could have argued that the only viable explanation of biological creativity is intelligent design. And it was true that before Darwin nobody had conceived of a viable alternative explanation. But that did not mean that Paley was right, of course. So, by the same logic, even if the Darwinian theory of human creativity is the only solution in sight at the moment, it does not follow that the Darwinian theory is correct.

What is especially challenging is to explain human hyper-creativity, i.e., the ability to create things that are *qualitatively different from, and more complex than,* previous creations. Stuart Hampshire recently voiced the way in which hyper-creativity drives both artistic and philosophical innovation.

> Compare the history of philosophy with the history of painting. A great painter appears, Giotto, Caravaggio or Cézanne, and nothing is ever the same again; his followers exploit the new opportunities opened up for them, and consequently we enter a new room in the museum. So also in the progress of philosophy, which draws upon the rare thinker of genius. (Hampshire 2002, p. 55)

Hampshire's point about artistic and philosophical innovation applies equally well to other areas of human creative endeavor. In each case, the same mechanism drives striking episodes of innovation. Great innovators (often called "geniuses") create new possibilities where none evidently existed before, and others follow their lead by exploring and exploiting the new opportunities made available.

It is relatively easy to imagine that Dennett's generate-and-test mechanism can explain the more mundane aspects of human creativity. Once a population of memes is alive in human culture, it is easy to imagine that variants of them could be produced and that natural selection would sift them for winners. But this is what could be called "mimetic" creativity rather than hyper-creativity. Mimetic creativity produces new variants on pre-existing themes, such as new fugues. Hyper-creativity produces qualitatively new themes that can subsequently become the fodder for mimetic creativity, such as producing the first fugue. It is harder to imagine that Dennett's generate-and-test mechanism can explain hyper-creativity. The key question is how to generate qualitatively new kinds of creations. The biological analogue of this is exactly what is missing in current constructive explanations of biological creativity. This gives us every reason to suspect that the mechanism behind human hyper-creativity will be similarly elusive.

In any event, my main point is that trading verbal theories back and forth is inconclusive. The most constructive approach to formulating and evaluating candidate explanations of human hyper-creativity is to produce a concrete system (probably computational) that exhibits some form of hyper-creativity. It might turn out that some sufficiently subtle generate-and-test mechanism will be sufficient to drive some such system. And it is plausible to suppose that generate-and-test mechanisms will play a significant role in any hyper-creative system. But some fundamentally new insight about a new kind of process might be another necessary ingredient in hyper-creative systems. The only way to answer these questions is to get one's hands dirty and actually build and experiment with concrete systems. One could summarize this methodological injunction with the slogan: "Put your model where your mouth is!" Artificial life is precisely the intellectual activity that is dedicated to following this slogan.

The positive picture tomorrow

Probably in the next decade artificial life will construct hyper creative systems. Achieving this goal will provide a distinctively generative light on both biological and human creativity. Along the way to this goal, artificial life will provide a positive and constructive methodology for intellectual progress on the issue.

Once artificial life has constructed hyper-creative systems and we have unlocked the key to biological and human hyper-creativity, society can exploit this new understanding. In essence, we will be able to create technology with all the creative and adaptive power and intelligence inherent in living systems. This new living technology will have myriad industrial and commercial applications. It will also be a fertile source of inspiration for a new generation of artists. For one thing, the ability to create living technology has far-reaching social and ethical implications. Artists are often inspired to participate in and to comment on social and ethical controversies. In addition, artists find new ways to augment their own creativity by harnessing the hyper-creative capacities of living technology. So, in the not-too-distant future we should expect artificial life to give birth to a new wave of hyper-creative living art.

References

Ando, Arike. 2001. What are People for? *Leonardo* 34, 447-452.

Bedau, Mark. 1996. The Extent to which Organisms Construct their Environments. *Adaptive Behavior* 4: 469-475.

Bedau, M. A. 1995. Three Illustrations of Artificial Life's Working Hypothesis. In Banzhaf, W., and Eeckman, F. (eds.) 1995. *Evolution and Biocomputation—Computational Models of Evolution* (Berlin: Springer), 53-68.

Bedau, M. A. 1996. The Nature of Life. In Boden, M., ed., *The Philosophy of Artificial Life* (New York: Oxford University Press), pp. 332-357.

Bedau, M. A. 2002a. Artificial Life. In L. Floridi, ed., *The Blackwell Guide to the Philosophy of Information and Computing* (Blackwell). In press.

Bedau, M. A., and Brown, C. T. 1997. Visualizing Evolutionary Activity of Genotypes. *Artificial Life* 5: 17-35.

Bedau, M. A., S. Joshi, and B. Lillie. 1999. Visualizing Waves of Evolutionary Activity of Alleles. In A. Wu, ed., *Proceedings of 1999 Genetic and Evolutionary Computation Conference Workshop Program* (Orlando: GECCO Proceedings), pp. 96-98.

Bedau, M. A., and Packard, N. 1992. Measurement of Evolutionary Activity, Teleology, and Life. In Langton, C., Taylor, C., Farmer, J. D., and Rasmussen, S., eds., *Artificial Life II* (Redwood City, CA: Addison Wesley), pp. 431-461.

Bedau, M. A., Snyder, E., Brown, C. T., and Packard, N. H. 1997. A Comparison of Evolutionary Activity in Artificial Systems and in the Biosphere. In Husbands, P. and Harvey, I., eds., *Proceedings of the Fourth European Conference on Artificial Life, ECAL97* (Cambridge: MIT Press), pp. 125-134.

Bedau, M. A., E. Snyder, N. H. Packard. 1998. A Classification of Long-Term Evolutionary Dynamics. In Adami, C., Belew, R., Kitano, H., and Taylor, C., eds., *Artificial Life VI* (Cambridge: MIT Press), pp. 228-237.

Bedau, Mark A., John S. McCaskill, Norman H. Packard, Steen Rasmussen, Chris Adami, David G. Green, Takashi Ikegami, Kunihiko Kaneko, Thomas S. Ray. 2000. Open Problems in Artificial life. *Artificial Life* 6: 363-376.

Bilotta, Eleonora, Pietro Pantano, and Valerio Talarico. 2000. Synthetic Harmonies: An Approach to Musical Semiosis by Means of Cellular Automata. In M. A. Bedau, J. S. McCaskill, N. H. Packard, and S. Rasmussen, eds., *Artificial Life VII* (Cambridge: MIT Press, 2000), pp. 537-546.

Crutchfield, J.P., Farmer, J.D., Packard, N.H., and Shaw, R. S. 1986. Chaos. *Scientific American* 255 (December): 46-57.

Dawkins, R. 19XX. *The selfish gene*. NewYork: Oxford University Press.

Dennett, Daniel. 1995. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*, Simon and Schuster: New York.

Dennett, Daniel C. 2001. In Darwin's wake, where am I? *Proceedings and Addresses of the American Philosophical Association* 75, pp. XX-XX.

Gould, Stephen Jay. 1989. *Wonderful Life: The Burges Shale and the Nature of History*, New York: Norton.

Hampshire, Stuart. 2002. The Spinoza solution. *New York Review of Books*, Oct 24, pp. 55-56.

Holland, J. H. 1992. *Adaptation in Natural and Artificial Systems*, 2nd edition, Cambridge: Bradford Books/MIT Press

Huws, Ursula. Nature, Technology and Art: The Emergence of a New Relationship? *Leonardo* 33, 33-40.

Innocent, Troy. 2001. The Language of Iconica. *Leonardo* 34, 255–260.

Langton, C. 1989. Artificial Life, in C. Langton, ed., *Artificial Life*, Redwood City: Addison-Wesley, pp. 1-47.

Langton, C. 1992. Life at the Edge of Chaos. In C. Langton, C. Taylor , D. Farmer, and S. Rasmussen, eds., *Artificial Life II*, Santa Fe Institute Studies in the Sciences of Complexity, Vol. X, Redwood City: Addison-Wesley, pp. 41-91.

Mignonneau, Laurent and Christa Sommerer. 2001. Creating Artificial Life for Interactive Art and Entertainment. *Leonardo* 34, 303 – 308.

Prophet, Jane. 2001. TechnoSphere: 'Real' Time, 'Artificial' Life. *Leonardo* 34, 309–312.

Rechtsteiner, A.and M. A. Bedau, 1999a, A Generic Model for Quantitative Comparison of Genotypic Evolutionary Activity. In D. Floreano, J.-D. Nicoud, F. Mondada, eds., *Advances in Artificial Life* (Heidelberg: Springer-Verlag), pp. 109-118.

Rechtsteiner, A. and M. A. Bedau, 1999b. A Generic Model for Measuring Excess Evolutionary Activity. In Banzhaf, W., Daida, J., Eiben, A. E., Garzon, M. H., Honavar, V., Jakiela, M., & Smith, R. E. (eds.), *GECCO-99: Proceedings of the Genetic and Evolutionary Computation Conference*, Vol. 2 (San Francisco, CA: Morgan Kaufmann), pp. 1366-1373.

Reynolds, C. W. 1987. Flocks, Herds, and Schools: A Distributed Behavioral Model. *Computer Graphics* 21, 25-34.

Reynolds, C. W. 1992. Boids Demos. In C. Langton, ed., *Artificial Life II Video Proceedings*, Redwood City: Addison-Wesley, pp. 15-19.
Ray, T. 1992. An Approach to the Synthesis of Life. In C. Langton, C. Taylor , D. Farmer, and S. Rasmussen, eds., *Artificial Life II*, Santa Fe Institute Studies in the Sciences of Complexity, Vol. X, Redwood City: Addison-Wesley, pp. 371-408.
Rumelhart, D. E., and McClelland, J. L. 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, 2 Vols., Cambridge: Bradford Books/MIT Press.
Skusa, A. and M.A. Bedau. 2002. Towards a comparison of evolutionary creativity in biological and cultural evolution. In R. Standish, H. Abbass, and M. A. Bedau, eds., *Artificial Life VIII*. Cambridge: MIT Press.Sober, E. 1994. The Adaptive Advantage of Learning versus A Priori Prejudice. In *From a Biological Point of View*, Cambridge: Cambridge University Press.
Sommerer, Christa, and Laurent Mignonneau. 2000. Modeling Emergence of Complexity: The Application of Complex System and Origin of Life Theory to Interactive Art on the Internet. In M. A. Bedau, J. S. McCaskill, N. H. Packard, and S. Rasmussen, eds., *Artificial Life VII* (Cambridge: MIT Press, 2000), pp. 547-554.
Whitelaw, Mitchell. 2001. The Abstract Organism: Towards a Prehistory for A-Life Art. *Leonardo* 34, 345–348.
Wolfram, S. 1994. *Cellular Automata and Complexity*. Reading: Addison-Wesley.
Wolfram, S. 2002. *A New Kind of Science*. Wolfram Publications.
Youngs, Amy M. 2000. The Fine Art of Creating Life. *Leonardo* 33, 377-380.