

How evolution guides complexity

Larry S. Yaeger¹

¹School of Informatics, Indiana University, 919 East 10th Street, Bloomington, Indiana 47408, USA

(Received 10 April 2009; accepted 29 August 2009; published online 19 October 2009)

Long-standing debates about the role of natural selection in the growth of biological complexity over geological time scales are difficult to resolve from the paleobiological record. Using an evolutionary model—a computational ecosystem subjected to natural selection—we investigate evolutionary trends in an information-theoretic measure of the complexity of the neural dynamics of artificial agents inhabiting the model. Our results suggest that evolution always guides complexity change, just not in a single direction. We also demonstrate that neural complexity correlates well with behavioral adaptation but only when complexity increases are achieved through natural selection (as opposed to increases generated randomly or optimized via a genetic algorithm). We conclude with a suggested research direction that might be able to use the artificial neural data generated in these experiments to determine which aspects of network structure give rise to evolutionarily meaningful neural complexity.

[DOI: 10.2976/1.3233712]

CORRESPONDENCE

Larry S. Yaeger:
larryy@indiana.edu

Although single-celled organisms may be every bit as well adapted to their ecological niches as human beings are to theirs, no one would argue that microorganisms are as complex as human beings, using any metric of complexity one might choose. Furthermore, looking at the fossil record it is clear that complexity, again by any metric, has increased over geological time scales (Carroll, 2001), from algae to plants, from ediacarans to arthropods to insects to mammals. Although this progression may not have taken the form of a simple ladder or ratchet (Gould, 1994), and due to horizontal gene transfer, endosymbiosis, and hybridization, our concept of the organization of branching species may have gone from tree to bush to thicket; we nonetheless have samples of a great many intermediate levels of complexity in the fossil record (and alive today) between microorganisms and human beings.

A natural question to ask is, how did this complexity arise? Does evolution itself foster a continuing escalation of organismal complexity? Will all evolutionary systems and all biologies produce such increases, thus producing intelligent life wherever life itself gains a foothold?

There have been many good arguments in support of the oft-held view that evolution is

biased toward increasing complexity. Rensch (1960a, 1960b) and Bonner (1988) argued that more parts will allow a greater division of labor among parts, and since this greater division of labor will confer an evolutionary advantage, evolution naturally favors a continuing increase in complexity as measured by a count of fundamental components. Waddington (1969) and Arthur (1994) suggested that due to increasing diversity, niches become more complex and are then filled with more complex organisms. Knoll and Bambach (2000) echoed and expanded on this argument concluding that evolution into an ever-increasing “ecospace” will confer a continual growth in complexity. Saunders and Ho (1976) and Katz (1987) suggested that component additions are more likely than deletions because additions are less likely to disrupt normal function, again resulting in an evolutionary increase in complexity as measured by parts. Kimura (1983), Huynen (1996), and Newman and Engelhardt (1998) demonstrated the value of neutral mutations in bridging gulfs in fitness landscapes through the selection for novel functions in previously neutral changes. These neutral mutations frequently take the form of gene duplications—a straightforward mechanism for increasing genetic complexity.

However, [Maynard Smith \(1970\)](#), [Raup *et al.* \(1973\)](#), [Gould \(1989, 1994, 1996\)](#), and others questioned whether the observed growth in complexity was the outcome of natural selection or simply, in Maynard Smith's words, the "obvious and uninteresting explanation" of a sort of random walk away from an immutable barrier of simplicity at the lower extreme—a growth in variance relative to the necessarily low complexity at the origin of life. Gould, in particular, argued extensively that random chance not only plays a greater role in evolution than previously understood but also is entirely sufficient to explain the observed increases in biological complexity over geological time scales. As part of an extended debate between them, [Dawkins \(1997\)](#) counter-argued that evolution must be progressive and biased even if that bias is not always in the same direction at the micro-evolutionary scale—in keeping with our experimental results reported here.

[Bedau *et al.* \(1997\)](#) and [Rechsteiner and Bedau \(1999\)](#) provided evidence of an increasing and accelerating "evolutionary activity" in biological systems that, until [Channon \(2001\)](#), were not observed in an artificial system. However, attempts to characterize complexity trends in the paleontological record have produced mixed results at best ([McShea, 1996](#); [Heylighen, 2000](#); [Carroll, 2001](#)), leaving us with no clear picture of the influence of natural selection on complexity. For more than a decade [McShea \(1994, 1996, 2001, 2005\)](#) has attempted to clarify and, where possible, empirically address the debate, by identifying distinct classes of complexity and, importantly, by distinguishing between "driven" trends in which evolution actively selects for complexity and "passive" trends in which increases in complexity are due simply to Gould's asymmetric random drift.

A number of researchers have used simple computational models of branching clade lineages to study driven vs passive trends in the evolution of complexity. [Raup *et al.* \(1973\)](#) looked exclusively at branching patterns. [Raup and Gould \(1974\)](#) used a ten-parameter vector to represent morphological characters. In [McShea \(1994\)](#), a single numerical parameter takes the place of complexity and biased or unbiased anagenetic (intraspecies) and cladogenetic (branching) random variation produces different populations. In all of these efforts, the distribution and statistics of the resulting populations are then at least roughly compared to the distribution and statistics of biological species over time in order to try to discern the presence or absence of any inherent bias in biological evolution. They largely conclude there is reason to doubt any systematic bias toward increasing complexity. Raup and Gould went so far as to suggest that "undirected selection may be the rule rather than the exception in nature." However, they base their arguments on the ability of random systems such as theirs to produce clade histories and temporal lengths of continuously changing characters that are merely similar to biological systems. They acknowledge that directional selection exists but emphasize that since the tem-

poral geometry of evolutionary phenomena can also be produced within purely random systems, such geometric patterns are insufficient to unequivocally imply directed causes.

McShea uses his model to conclude that a biased system will necessarily exhibit an increase in the minimum complexity present in the system and this observation has become common wisdom ([Wagner, 1996](#); [McShea, 2001](#); [Carroll, 2001](#)). The presence of a multitude of minimum complexity, single-celled organisms today is then taken as evidence against a biased evolutionary trend in complexity. However, even though an increase in the minimum is evidenced in the most commonly reported version of McShea's model, he also notes that it is not always possible to use this observation to distinguish a passive system from a weakly driven one. Moreover, in a more realistic evolutionary system in which fitness at smaller, less complex scales is relatively independent of fitness at larger, more complex scales, the system could possess a substantial positive bias at larger scales without eliminating or even disadvantaging organisms at the lower end of the spectrum. Indeed, in the real biological realm, the largest, most complex organisms are themselves extremely hospitable niches for single-celled organisms and, as opposed to competing with microorganisms in any way, significantly increase their available resources.

One aspect of McShea's model that tends to color its predictions, besides its sheer simplicity, is its use of an anagenetic change parameter. This within-lineage change increases the "complexity" of the lineage without producing a new branch. While anagenetic change may be observable in nature, fine-tuning some morphological or behavioral character of a species, we think it likely that nontrivial changes in complexity will more commonly be associated with speciation events suggesting that anagenetic changes in the model should be of a substantially smaller magnitude than cladogenetic changes, which is not generally the case in McShea's experiments. Additionally, this kind of anagenetic change is functionally equivalent to a cladogenetic branch followed by an immediate extinction event of the ancestor lineage. With McShea's nominal set of model parameters, it is as if half of all speciation events resulted in a fierce competition for the same resources and the parent lineage always lost. But speciation is likely to be associated with adaptation to a different niche, so again this anagenetic change component of the model is likely to produce an unrealistic trend in population statistics. [McShea \(1994\)](#) acknowledged a dramatically lower, nearly unmeasurable rate of growth in the observed minimum in a purely cladogenetic model compared to a mixed anagenetic and cladogenetic model. It is therefore clear that reducing or eliminating this anagenetic component will significantly reduce the observed growth in the minimum for positively biased systems.

Another aspect of McShea's model that is likely to skew the results relative to a more realistic model of evolution is his choice of extinction rates. At first glance it is difficult to

fault his simple, seemingly unbiased choice—a small, constant extinction rate that applies equally to all lineages. But particularly since his increasing numerical parameter is supposed to represent complexity, it may well be important to recognize that complex things break down more easily than simple things. [Farmer and Griffith \(2007\)](#) studied the robustness of self-reproducing machines and derived an intriguing functional relationship between viability and complexity that might improve McShea's model. An even greater effect on extinction rates might be due to population size. A small population is obviously more easily extinguished than a large population. Accounting for actual population differences between microorganisms and mammals, say, while difficult to even estimate, would almost certainly produce an essentially negligible extinction rate for the former relative to the latter. Simply accounting for reasonable growth rates relative to lineage origin times is likely to yield a substantial difference in extinction rates between later, more complex lineages and earlier, simpler lineages. We would expect the modification of extinction rates according to these viability and population principles in a McShea-like model to substantially reduce any growth in the minimum for driven systems.

Given the limitations of these simple models and the difficulties and ambiguities one finds when studying the paleontological record, it makes sense to turn to computer models that actually employ evolution to investigate the question of evolutionary trends in complexity. [Turney \(1999, 2000\)](#) used a simple evolutionary model to suggest that increasing evolvability is central to progress in evolution and predicts an accelerating increase in biological systems that might correlate with complexity growth. [Adami *et al.* \(2000\)](#) and [Adami \(2002\)](#) defined complexity as the information that an organism's genome encodes about its environment and calculates a complexity metric that is inversely proportional to the entropy of its genetic elements (bits in simulation, base pairs in biological genomes). His AVIDA software has been used to show that asexual agents in a fixed, single niche always evolve toward greater complexity of this uniquely defined type. Although we believe there is little doubt that Adami's observation regarding genes encoding information about the environment is both true and important, we suspect that his "complexity" metric is better viewed as a measure of a species' adaptation to an environment or, more simply and directly, a measure of genetic consistency due to its failure to account for multiple species, multiple and variable niches, niche creation, or any kind of biodiversity.

In previous work ([Yaeger and Sporns, 2006](#)), we have used our computational ecosystem software, Polyworld, together with an information-theoretic measure of neural complexity to demonstrate that complexity of this particular type does increase over evolutionary time scales. That earlier work, however, did not distinguish contributions from natural selection versus contributions from Gould's random walk. Basically, there was no "null model" to compare with in that

early work and since the seed population used to initialize those simulations consisted of very low-complexity agents, we could not rule out Maynard Smith's obvious and uninteresting explanation. Subsequently ([Yaeger *et al.*, 2008](#)), we established a mechanism for running a kind of null model in Polyworld and were able to tease out the differing contributions of natural selection and random drift, and thus demonstrated a statistically significant selection for complexity growth under certain circumstances. However, perhaps predictably, the story was not as simple as "evolution selects for complexity growth." Rather, it proved to be the case that evolution, as Dawkins argues, always drives complexity growth (or loss or stabilization) in a biased fashion; however, it is also the case that evolution does not always drive complexity in the same direction, as will be discussed below. We also show that periods of complexity growth correspond to periods of behavioral adaptation, consistent with the findings of [Seth and Edelman \(2004\)](#) suggesting our neural complexity metric is a good measure of what might be considered primitive and adaptive forms of intelligence. However, we further demonstrate that not all forms of neural complexity confer an evolutionary advantage suggesting that in order to relate artificial neural network models to biological brain networks it may be necessary to study networks specifically evolved for their behavioral adaptations in an ecological context subject to natural selection.

Note that variants on the word "guide" take on more than one meaning herein. Like most of the papers in this special issue on "guided self-organization," we sometimes refer to human intervention and selection of outcomes in otherwise self-organizing processes. This is particularly the sense when we discuss the use of complexity as a fitness function expressly guiding evolution. We also conclude with remarks on future research directions that include possible uses of this work to further hone and guide evolution's outcomes in our artificial life models. But throughout most of the paper we are addressing the degree to which natural selection, itself the overarching self-organizing process for all of biology, "guides" the evolution of organismal complexity, in a biased (or unbiased) fashion. With the distinction made, usage should be clear from context.

SIMULATION SOFTWARE

The software model used for this study is Polyworld ([Yaeger, 1994](#)), an evolutionary model of a computational ecosystem, populated by haploid agents with a suite of primitive, neurally controlled behaviors (move, turn, eat, mate, attack, light, and focus). The artificial neural networks (ANNs) that control these agents use "summing and squashing" (a.k.a. "firing rate") neurons and perform Hebbian learning at the synapses. The architecture of the ANN is derived from the agent's genome, expressed as a number of neural groups of arbitrary excitatory and inhibitory neuron counts, along with connection densities, topologies, and learning rates between




Figure 1. The Polyworld evolutionary ecosystem. The multicolored trapezoids are the agents (with color denoting certain behaviors and genetics): bright green blocks are food; the brown walls are the barriers the agents cannot pass through. All actions take place on the dark green ground plane. Also shown along the top are certain graphical statistics (left), the point of view of all agents in the world (left-center), a brain connection matrix (center), an overhead detail view (right-center), and a few numerical statistics (right).

these neural groups and types. Input to the ANN consists of a row of pixels from a rendering of the scene from each agent’s point of view, like light falling on a retina. Although agent morphologies are simple and static, agents interact with the world and each other in fairly complex ways. They reproduce through the simultaneous expression of a mating behavior by two collocated agents.

The agents’ energy stores are normally depleted by all actions, including neural activity with stronger actions (higher behavioral neuron activations) and larger neuron and synapse counts depleting more energy. To simplify the analysis the per-neuron and per-synapse neural costs have been eliminated for this series of experiments; all behavior-based costs have been retained. Energy is also depleted when one agent is attacked by another. Energy must be replenished by seeking out and consuming food in the environment or by killing and eating other agents. Minimum and maximum population sizes are controlled by proportionately decreasing energy costs as the population dwindles and increasing energy costs as the population swells. In its normal mode of operation all evolutionary pressures emerge from a process of natural selection—the survival and reproduction of the agents—and there is no influence of any kind from a fitness function (although an *ad hoc* heuristic fitness function is computed for purely informational purposes).

Polyworld was originally designed to explore the possibility of open-ended evolution of artificially intelligent agents. Figure 1 shows the Polyworld evolutionary ecosystem and various attendant graphical elements. For further background on the simulation software please see [Yaeger \(1994\)](#).

COMPLEXITY

To be useful scientifically, complexity must be quantifiable. Complexity for the purposes of evolutionary biology has taken many definitions and forms over the years, from organism size ([Cope, 1871](#)) to distinct cell types ([Bonner, 1988](#); [Valentine et al., 1994](#)) to morphology ([Thomas and Reif, 1993](#); [McShea, 1993](#)) to ecological webs of interaction ([Knoll and Bambach, 2000](#)). Information theory ([Shannon, 1948](#)) has produced various methods for defining and quantifying complexity that might be suitable for the study of biological and evolutionary processes but most such measures share an undesirable trait with Shannon entropy, which is that they are maximized by random systems in which all states are equiprobable. By contrast, biological complexity corresponds to neither minimal nor maximal entropy but something in between. Although unaware of his work at the time, the inventors of the complexity measure used here shared an insight with [Chaitin \(1979\)](#): the way information flows

across multiple scales relates directly to the internal structure of a system. Chaitin also defined a “measure of mutual information for n -tuples” that presages a key component of our complexity measure (integration, a.k.a multi-information); however he stopped short of applying these insights to his algorithmic entropy measure of complexity.

Tononi *et al.* (1994) proposed an information-theoretic measure for quantifying the complexity of neural dynamics, applicable to arbitrary sets of random variables, which avoids maximization by randomness, instead capturing quantitatively the manner in which information is processed at multiple scales. Widely acknowledged as an important conceptual milestone in the measurement and understanding of complexity, the Tononi–Sporns–Edelman measure is now commonly referred to as “TSE complexity” (for the authors’ initials). A simpler, much more computationally efficient measure was subsequently proposed in Tononi *et al.* (1998) and explored computationally in Sporns *et al.* (2000). The methods were subsequently refined in Lungarella *et al.* (2005) and released as a MATLAB Complexity Toolbox, currently maintained at http://www.indiana.edu/~cortex/complexity_toolbox.html/. In order to be able to process the neural dynamics from hundreds of neurons for thousands of agents, per each of the many simulations, the simplified TSE complexity measure was re-implemented in C++ and is what we use and refer to throughout the remainder of the paper as “complexity” or C .

Although nontrivial to derive and implement, the intuition behind the original TSE complexity is straightforward: cooperation among various elements of a network, called integration and measured by a multivariate extension to mutual information, increases network complexity, to a point. But specialization of network subunits, called segregation, also increases network complexity. Segregation produces a difference between an estimate of integration that is uniformly distributed across all scales (subsets of units) and the actual measured integration at these different scales. This difference is what TSE define as complexity. Maximal complexity is thus achieved in networks that maximize both integration and segregation—both cooperation and specialization—to the extent possible. The original TSE complexity is given by

$$C_N(X) = \sum_{k=1}^n \left[\langle H(X_j^k) \rangle - \frac{k}{n} H(X) \right], \quad (1)$$

where $H(X)$ is the Shannon entropy of the entire system of n variables, k is the size of a subset of variables, and j indicates that the ensemble average $\langle H(X_j^k) \rangle$ is to be taken over all $n! / (k!(n-k)!)$ combinations of k variables. The simplified approximation looks only at the first subset term, involving mutual information between individual variables and the rest of the system, or cast in terms of entropies

$$C(X) = H(X) - \sum_{x_i \in X} H(x_i | X - x_i), \quad (2)$$

where $H(X)$ is again the entropy of the entire system and the $H(x_i | X - x_i)$ terms are the conditional entropy of each of the variables x_i given the state of the rest of the system.

NATURAL SELECTION (DRIVEN) VS RANDOM DRIFT (PASSIVE)

Although it would be difficult to eliminate natural selection from an evolving biological ecosystem, artificial ecosystems such as Polyworld are more flexible. In order to distinguish between driven trends in complexity due to natural selection vs passive trends due to random genetic drift, we designed a “lockstep” mode of operation for the simulator in which a passive run, where natural selection is disabled, is tied to a driven run, where natural selection operates like normal. First the normal, driven simulation is carried out during which all agent births and deaths are recorded (along with all relevant neural data). Then the lockstep, passive simulation is carried out in which agents are not allowed to reproduce or die as a result of their own behaviors. Instead, every time an agent died in the driven simulation, an arbitrary agent in the passive run is chosen at random and killed. Similarly, every time two agents reproduced and produced an offspring in the original, driven simulation, two arbitrary agents in the passive run are bred at random and the resulting offspring is placed at a random location in the world.

By this method, population statistics between the driven and passive runs are identical. As a result, the statistics of all genetic operations—crossovers and mutations—are comparable between the paired runs. The number of crossover points and mutation rates are themselves embedded in the genome (Yaeger, 1994), so the genetic operations are only the same statistically, not in detail. Similarly, the “life experiences” of a given agent—its trajectory through the world and the inputs to its visual system—are only comparable statistically between paired runs. Since the agents’ life experiences do have an impact on the values of neural complexity we compute, this may produce some differences between paired runs but we do not expect this to have a large influence, due to a combination of statistically similar environments (and therefore similar visual inputs) and only modestly different observed agent trajectories in these paired runs. Due to the resulting consistency in visual environment, the controlling statistics, such as the entropy and mutual information in the visual inputs, should be comparable between paired runs, thus producing comparable complexity values. Regardless, specifically because complexity is affected by agent behaviors and their resulting sensory inputs, one could argue that agents in lockstep runs must be allowed to control their own actions in order to obtain valid measures of their neural complexity. So even though we expect observed complexity differences to be primarily the result of

differences in neural architecture, any differences due to agent behaviors and resulting life experiences are appropriately included in the comparisons.

The effective result is that natural selection is “turned on” in the driven runs and “turned off” in the passive runs. For example, gene states are subject to natural selection, based on the evolutionary viability—the fitness—of the agents’ behaviors in the driven runs, while gene states are subject only to the same degree of variation, with no evolutionary fitness consequences or effects, in the passive runs. The passive run serves as our null model, thus allowing us to distinguish between complexity growth due to natural selection vs that due to a random walk in gene space.

Note that while genetic changes in both the driven and passive cases produce different neural architectures, neural dynamics, and agent behaviors, in the passive case (only) these changes have absolutely no impact on the reproductive success of agents. Thus this random walk in gene space corresponds to and produces a random walk in complexity space (as measured by the neural dynamics), without regard for the evolutionary value of that complexity or the corresponding agent behaviors. By contrast, in the driven case the variations in neural structure and function and resulting behavioral adaptations produced by these genetic changes directly affect the agents’ reproductive success. Accordingly, there is no expectation of any strict correlation between the agents’ behaviors in the driven and passive runs. Especially, there is no expectation that agents in the passive runs will exhibit the kinds of foraging and mating behaviors seen in driven runs.

DATA GENERATION AND ACQUISITION

A simple world was set up in Polyworld that accommodated between 90 and 300 agents, grew food uniformly and randomly in two patches at opposite ends of the world (80% at one end and 20% at the other), positioned two barriers that ran 90% of the depth of the world, and divided the world into thirds across its width (see Fig. 1). The agents were allowed to evolve between zero and five internal neural groups with up to 16 excitatory and 16 inhibitory neurons per group. Input neural groups, internal neural groups, and output/behavior neural groups were connected according to evolved connection densities specified in the agents’ genomes. This allows up to 217 neurons and 45,584 synapses. Each simulation was allowed to run for 30,000 time steps, which corresponds to about 400 generations. For each run, the neural activation of every neuron, at every time step, for every agent is recorded to disk, yielding approximately 10 GB of data. It is this data to which we subsequently apply our complexity calculations, thus allowing us to determine the neural complexity of each individual agent. Note that an agent’s complexity is only fully determinable upon its death—when its neural activation time series is complete. We therefore exam-

ine mean trends in the population’s complexity for all agents that have died during a particular time interval—typically every 1,000 time steps.

Ten driven runs were carried out, in which natural selection operated normally, varying only the initial seed to the random number generator. A corresponding passive “lockstep” run, in which natural selection was disabled, was carried out for each of the driven runs yielding 20 simulations in all.

The runs were all seeded with a population of agents based on a uniform genome. This seed genome was designed to produce modest dispositions toward potentially useful behaviors, such as running toward green (food) and away from red (attacking agents); see [Yaeger \(1994\)](#) for details on color use in Polyworld. However, these seed agents are not a viable species. That is, test runs have demonstrated that without the ability to evolve the seed agents will fail to survive and reproduce in numbers sufficient to sustain their population, as a result of which their population will dwindle and go extinct. There must therefore be a significant evolutionary value to some kind of rewiring of the genome and resulting neural network architectures as a result of natural selection, at least until such point as the agents’ behaviors are sufficiently adapted to the environment to sustain their numbers through their foraging and mating behaviors. The genome used in this seed population produces neural architectures that are nearly minimal in size and complexity, so one should expect the random variations in gene space associated with the passive, lockstep runs to also yield increases in complexity, as indeed they do. It is primarily the differences in the rate of change in complexity between driven and passive runs that are used to assess the role of natural selection in complexity increase.

Complexity may be calculated for various subsets of neurons—all neurons, just the input neurons, just the behavior neurons, or the “processing” neurons (all neurons except inputs). All complexities presented here are based on the processing neurons, although, in general, there are few differences in complexity trends between all neurons and processing neurons. Complexity then varies as evolution produces changes in the parts of the genome that specify the neural architecture and thus alters the internal neural dynamics and resulting behaviors of the agents.

COMPLEXITY AS A FITNESS FUNCTION

Given a system such as Polyworld that is capable of evolving a wide range of neural architectures and using them to control a fairly broad range of agent behaviors, and given the ability to then calculate an information-theoretic measure of complexity for these agents’ neural networks, an intriguing possibility presents itself: why not dispense with natural selection altogether and instead directly employ neural complexity as a fitness function, in the traditional genetic algorithm (GA) sense? To investigate the fruitfulness of such an




Figure 2. Driven and passive complexity vs time. Light solid lines show population-mean complexity for individually driven, natural selection runs. Light dashed lines show population-mean complexity for individual passive, lockstep runs. Heavy lines show means of all ten runs for corresponding line style. Light dotted line at bottom shows $1-p$ -value for a dependent student's T -test with a horizontal $p > 0.05$ T -critical dashed line at 0.95.

approach, in addition to the paired driven vs passive runs just described, we carried out a small number of simulations using neural complexity as a fitness function. This required us to run Polyworld in a unique, third (neither normal nor lockstep) simulation mode in which it acts as a steady-state GA. In this mode the simulator retains a list of the N (typically 30) best agents that have ever lived, according to a specified fitness function. For these runs we computed neural complexity of the processing units for each agent upon its death and used this complexity directly as a measure of the agent's fitness. Then, only in this steady-state GA mode, upon the death of each agent the software takes an incremental step through the list of N best agents, using a pair of indexes, one of which only increments when the other index finishes its run and wraps around, and uses the two agents identified by these indexes to produce an offspring, employing the usual mechanisms of crossover and mutation. (A simple experiment using a fitness function of $1/(\text{velocity} + 1)$, in which agents were successfully evolved to maintain near zero velocity population-wide, was used to validate this technique.) Results of these simulations will be discussed below.

RESULTS AND DISCUSSION

The resulting population-mean complexities for all 20 simulations are shown with light lines (solid for driven, dashed

for passive) in Fig. 2, along with additional, bold lines corresponding to the means of the two classes of runs (driven vs passive). In addition to the simulation results, a measure of statistical significance is shown at the bottom of the graph. The fine dotted line shows $1-p$ -value for the dependent student's T -test. The horizontal dashed line at 0.95 thus corresponds to a $p > 0.05$ critical threshold. Where the dotted line rises above the dashed line, the difference between the driven and passive runs may be considered statistically significant; where it falls below, the difference is not significant. Where the driven and passive mean complexity values cross ($t=7,000$) significance drops almost to chance (0.5), as one would expect, but displayed values have been truncated below to 0.8 for clarity.

The first thing to note is that during the early stages of the simulation evolution is actively selecting for increasing complexity. This is consistent with our observation that the seed population is too simple to sustain its numbers and must evolve or become extinct. As long as complexity is of an evolutionary advantage, natural selection will behave in this biased, driven fashion with regard to complexity.

However, note that any statistical significance has disappeared by about $t=4,000$ and indeed, the mean passive complexity reaches the same level as the mean driven complexity by

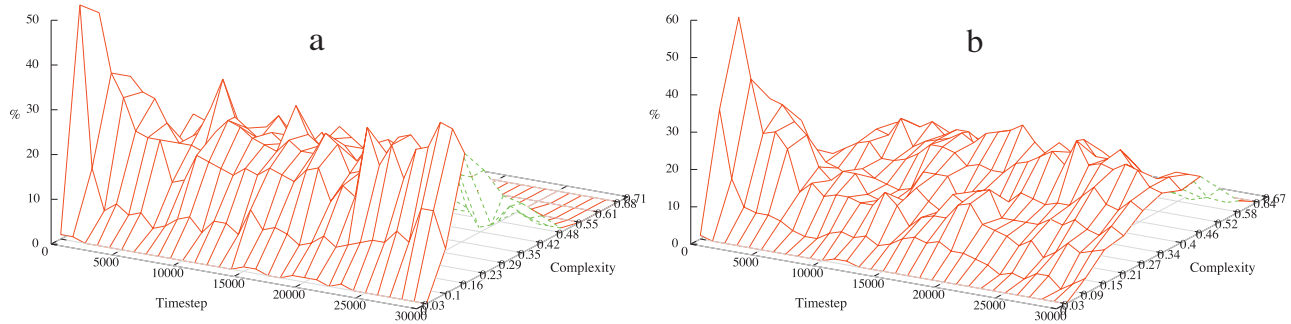


Figure 3. Population histograms of complexity over time for a driven run (a) and a passive run (b).

about $t=7,000$. At roughly the point that statistically significant differences disappear, the driven complexities produced by natural selection begin to plateau. What is going on is that by this stage, a “good enough” solution has emerged and begun to spread throughout the population, as evidenced by the peakedness of the temporal histogram of complexity in Fig. 3(a). From this time forward, in the majority of the simulations, natural selection tends to act as a weakly stabilizing force, maintaining a nearly constant mean level of fitness in the population. Since most genetic variation at this point is likely to be of a deleterious nature, evolution acts to weakly maintain this good enough solution. Meanwhile, the random walk continues in the passive runs, thus increasing variance in the genes and in complexity, as evidenced by the flat, broadened complexity distributions of Fig. 3(b). The growing variance produces networks of greater and greater complexity—now significantly exceeding that of the driven runs, on average.

The stabilization effect consistently observed in driven runs is, however, only a weak effect, as attention to the individual driven runs—the light, solid lines—reveals that in several of the runs, in a kind of punctuated equilibrium, variation eventually produces a network that is sufficiently better adapted to the environment that its genes spread throughout the population causing the mean level of complexity in the population to rise again. These new driven solutions reach approximately the level of the passive complexity and then plateau once again at this new, higher level. With more evolutionary time it is possible that more of the driven runs would make at least this secondary leap in complexity, bringing the overall mean driven level of complexity up to the level of complexity seen in the passive runs or higher.

The overall mean passive complexity also plateaus in these runs but for a much less interesting reason. By about $t=10,000$ the bits of the agent genomes have been fully randomized to roughly 50% on and 50% off. At this point, variance is at a maximum in the model and further randomization merely shuffles bits around without any change in the mean values of the genes that control the topology of the neural networks. It is due to the fixed size genome and fundamentally linear interpretation of genes in Polyworld that such an upper bound on variance exists and once maximum variance is reached the passive

runs cease to characterize a truly random walk. Complexity would, of course, continue to increase in an unbounded random walk, which would only increase any statistical significance of differences between passive and stabilized driven trends observed at these later times.

Note that final complexities range from about 0.25 to 0.35 and contrast this to the 0.9 level of complexity achieved by a steady-state GA simulation using complexity as its fitness function, as shown in Fig. 4. The factor of three multiplier indicates that this method for evolving larger values of complexity certainly works. However, the behaviors of the agents in this complexity-as-fitness run are not well suited to the normal evolutionary constraints of life in Polyworld. The agents do not forage for food. There is no indication that they seek out other agents for reproduction. This run was seeded with agents based on the same seed genome as the driven and passive runs. So initially they are not particularly well suited to their environment. After about 6,000 or so time steps, at a moderate level of complexity (around 0.75), most of the agents have adopted a stereotypical tight turning behavior, all spinning in small loops, and ignoring food and other

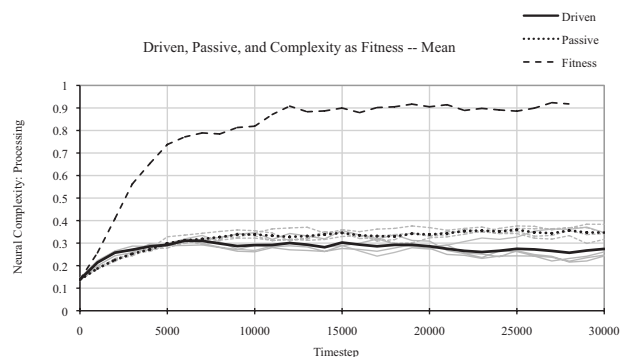


Figure 4. Complexity from several driven runs vs passive runs plus a single complexity-as-fitness run. Light solid lines show population-mean complexity for individual driven runs. Light dashed lines show population-mean complexity for individual passive runs. Heavy solid and dashed lines show means of the individual runs for the corresponding line style. Heavy line with long dashes shows result from a single run using complexity as a fitness function in a steady-state GA mode of operation.

agents. We note that this behavior produces fairly high entropy while maintaining a significant degree of mutual information in the sensory units (due to adjacent pixels being correlated and some degree of frame-to-frame coherency), thus producing a fairly high degree of complexity throughout the network. Later, near peak complexity the agents' behaviors have become more interesting and more varied (although some tight-looping agents remain) but they still do not explicitly seek food, and there is still no evidence of any substantial degree of agent-agent interactions (such as would be required for reproduction).

The driven vs passive simulation results clearly demonstrate that in this model system natural selection will drive increases in complexity, in a biased fashion. Furthermore, it is not unreasonable to expect results from the model to apply to biological systems, since, even though Polyworld provides only an abstract model of biological evolution, the principle in action here is, at its base, simply natural selection—an amplification of advantageous traits. The mechanisms of action and behavior obviously differ between model and biology, and there will be differences in the specific costs and benefits between biological agents and artificial agents (just as there are between one biological agent and another). However, there will always be cost/benefit tradeoffs in competition and niche exploitation for any evolutionary ecosystem, real or artificial, and it is these tradeoffs that evolution explores. As long as increases in complexity improve the ability to survive and reproduce by agents so advantaged, those complexity increases will be actively selected for. The subsequent plateau in driven complexity seen here, due to the widespread adoption of a good-enough strategy that acts as a kind of local optimum, demonstrates how evolution can act to serve as a weak stabilizing force for complexity, thus echoing and providing a concrete example of Dennett's (Dennett 1996) observation, "The cheapest, least intensively designed system will be 'discovered' first by Mother Nature, and myopically selected."

Yet when variation produces a sufficiently improved network and corresponding set of behaviors, natural selection draws a punctuation mark on the old era and ushers in a new equilibrium, as so often observed in nature.

Furthermore, the early period during which the growth rate of complexity in the driven runs exceeds that in the passive runs corresponds to a period of behavioral adaptation of the agents. Prior work (Griffith and Yaeger, 2006) demonstrated that Polyworld agents evolve a near ideal free distribution (IFD) of agents to resources (Fretwell and Lucas, 1970; Fretwell, 1972), for a variety of different distributions of food in patches. In these simulations agents have evolved to mostly inhabit the food patches, in IFD proportions by approximately $t=4,000$ and exhibit little change in food patch occupation beyond about $t=7,000$ (see Fig. 5). These times correspond to the period when driven complexity reaches its peak and plateaus and after which passive complexity surpasses

driven complexity. The ecosystem simulated here is deliberately simple and it is sufficient for the agents inhabiting it to forage and find mates efficiently, which conditions are satisfied by simply occupying these fixed food patches and frequently expressing their eating and mating behaviors. This behavioral adaptation is reflected in the approach to IFD conditions and corresponds directly to the period of increasing neural complexity when natural selection drives genetic change. When genetic change and corresponding changes in neural complexity are random, however, as in the passive runs, agents (predictably) never adopt an IFD, thus demonstrating that our neural complexity metric and behavioral adaptation are not intimately bound when natural selection is absent.

The complexity-as-fitness-function results also demonstrate that this measure of neural complexity is not uniquely tied to behavioral adaptation and evolutionary fitness—to the kinds of survival and reproduction that an evolutionary biologist would construe as fitness. So complexity of a behavioral network's dynamics alone is not sufficient to guarantee survival of an agent, or selection by evolutionary pressure. Although this seems at odds with the much more direct coupling of behavior and complexity found in a task discussed by Sporns and Lungarella (2006) in which a robotic arm attempts to grasp a randomly moving block, the discrepancy is almost certainly due to significant differences in the range of possible behaviors, environments, and sensory inputs in these two experiments. In the work by Sporns and Lungarella, evolving the robot for maximum complexity produced block-grasping behaviors as efficient as when the robot was evolved to perform the block-grasping task directly. But in their experiment, aside from the block, the environment was empty or consisted of random noise and the only way to provide nonrandom sensory input to the controlling network, and thus increase the complexity metric, was to attend to that block. By contrast, in Polyworld, the environment contains appearing (growing) and disappearing (eaten) food, other structural artifacts (the barriers), and most significantly, many other agents behaving according to their own network dynamics, providing a rich and varied source of visual complexity. Each Polyworld agent also expresses control over a suite of seven primitive behaviors, most of which directly affect that agent's experience of the world. All these additional sensory inputs and multiple low-level behaviors produce a complicated, nonunique relationship between the complexity of the agent's neural dynamics and any resulting higher order behaviors. It is the extreme difference between the agents' sensoria and possible sources of complexity between this work and that of Sporns and Lungarella that we believe accounts for the marked difference in correlation between complexity and behavior in the two studies.

While the driven results show that TSE complexity is well correlated with the evolution of adaptive behaviors in an ecosystem subject to natural selection, the fact that complexity in the passive case ultimately surpasses that in the driven

Food distributed in two patches 80% / 20%

Distribution of Agents to Food Resources

Driven ——— Passive - - - - -




Figure 5. Distribution of agents to heterogeneous resources. In the driven case (solid lines) agents approach an IFD over the same time period that evolution actively selects for complexity growth. In the passive case (dashed lines) they do not.

case demonstrates that such a correlation does not exist when natural selection is absent. This is consistent with the complexity-as-fitness-function results, where, again, complexity growth in the absence of natural selection is shown to be uncorrelated with evolutionarily useful behavioral adaptation.

CONCLUSIONS

Three key components—a computational ecosystem capable of carrying out synchronized natural-selection and null-model simulations, a resulting wide array of evolved neural network architectures, and an information-theoretic metric for capturing the complexity of the neural dynamics in these networks—has allowed us to ask questions and observe trends that are difficult to pose and assess by any other means. By “replaying the tape of life” (Gould, 1989) in these novel ways, we have been able to investigate evolutionary trends in complexity, a subject much discussed and debated since Darwin. Our results demonstrate a clear natural selection for complexity in a driven, biased fashion. But they also show a tendency to weakly stabilize complexity at a “just good enough” level. Although our work does not specifically address it, there is little doubt that evolution will, under the right circumstances, select in a driven sense against complexity, such as when dark-dwelling organisms in a cave give

up their eyes to avoid wasting energy on unneeded complexity. Thus it can be seen that, at the scale of individuals comprising a species, evolution always guides trends in complexity.

But the scale of the discussion matters a great deal. Gould (1996) and Dawkins (1997) argued strongly for passive and driven evolutionary trends, respectively. However, much of their disagreement seems to stem from an issue of scale. Our work agrees with Dawkins’s claim that evolution is always driven—at the level of individuals comprising a species. But integrating over an entire biome’s collection of increasing, decreasing, and stabilized trends in complexity is likely to produce a process that at least appears random, at the larger scale. In one of the earliest attempts to model evolution computationally in order to address these kinds of questions, Raup *et al.* (1973) observed that fully deterministic trends at small scales may very well be at the base of larger scale trends, even if those larger scale trends turn out to be passive. Our current simulations reinforce this notion of natural selection driving the evolution of complexity at small scales, but driving it in all directions—up, down, and stable. So larger scale trends are likely to be obfuscated by these opposing trends at the smaller scale, and the often conflicting and inconclusive evidence for trends in the paleobiological record is unavoidable and to be expected.

Finally, by looking at a specific adaptation—the evolution toward behaviors resulting in an ideal free distribution—we were able to demonstrate a correlation between our measure of neural complexity and behavioral adaptation. But results from the driven vs passive study and the complexity-as-fitness-function study make it clear that this correlation only applies when complexity gains are the result of natural selection.

FUTURE DIRECTIONS

At this stage it is only conjecture, but we suspect that, despite Gould's arguments, the evolution of complexity will be seen to behave in a driven fashion at larger scales as well, and that we have some hope of addressing this larger scale issue in our model. If Waddington (1969), Arthur (1994), and Knoll and Bambach (2000) were correct about complexity naturally increasing due to an ever-expanding ecospace—a continuing growth, production, and occupation of new, more complex evolutionary niches—then increases in the range, complexity, and ease of production of new niches in Polyworld should be matched by a concomitant increase in the complexity of the neural dynamics of the agents evolved to occupy those niches. We hope there will be ethological level, qualitative evidence for these complexity increases, but regardless, our information-theoretic complexity measure will allow us to quantitatively assess any such trends.

Another area we hope to address in future research is a principled assessment of the relationship between network structure and function. Our process generates a substantial record of network architectures and corresponding network dynamics. The burgeoning field of Network Science has demonstrated the value of applying graph theoretical methods to the understanding of a wide variety of natural and artificial systems. There is already evidence for a consistent relationship between distributions of network motifs in biological brains and complex artificial neural networks (Sporns and Kötter, 2004; Alon, 2007). It may be possible to identify network structural metrics in our data that are predictive of complex network function. Initial steps in this direction have shown that as complexity increases in Polyworld nets, the clustering coefficient increases and the average minimum path length decreases, suggesting an evolution toward “small world” network characteristics (Lizier *et al.*, 2009). Interestingly, given the observed disconnect between complexity and behavioral adaptation when natural selection is absent, it may even be important to carry out this kind of analysis on networks evolved for their behavioral, evolutionary effectiveness, rather than for arbitrarily complex dynamics. Finally, if successful in characterizing those network structures most likely to confer dynamical complexity, we may be able to close the loop on guided self-organization and apply our learned structural constraints to the more rapid, more far-reaching evolution of artificially intelligent agents.

ACKNOWLEDGMENTS

Much of the work discussed in this paper was carried out with Olaf Sporns (Psychological and Brain Sciences, Indiana University) and Virgil Griffith (Informatics, Indiana University and Koch Computational Neuroscience Laboratory, California Institute of Technology). My grateful thanks to the three anonymous reviewers whose comments significantly improved the content of this paper.

REFERENCES

- Adami, C (2002). “What is complexity?” *BioEssays* **24**, 1085–1094.
- Adami, C, Ofria, C, and Collier, T (2000). “Evolution of biological complexity.” *Proc. Natl. Acad. Sci. U.S.A.* **97**(9), 4463–4468.
- Alon, U (2007). “Network motifs: theory and experimental approaches.” *Nat. Rev. Genet.* **8**, 450–461.
- Arthur, B (1994). “On the evolution of complexity.” *Complexity: Metaphors, Models, and Reality*, Cowan, GA (ed), pp. 65–81, Addison-Wesley, Reading, MA.
- Bedau, M, Snyder, E, Brown, C, and Packard, N (1997). “A comparison of evolutionary activity in artificial evolving systems and in the biosphere.” *Proceedings of the Fourth European Conference on Artificial Life*, Husbands, P, and Harvey, I (eds), pp. 125–134, MIT Press, Cambridge, MA.
- Bonner, J (1988). *The Evolution of Complexity by Means of Natural Selection*, Princeton University Press, Princeton, NJ.
- Carroll, S (2001). “Chance and necessity: the evolution of morphological complexity and diversity.” *Nature (London)* **409**, 1102–1109.
- Chaitin, GJ (1979). “Toward a mathematical definition of ‘life’.” *The Maximum Entropy Formalism*, Levine, RD, and Tribus, M (eds), pp. 477–498, MIT Press, Cambridge, MA.
- Channon, AD (2001). “Passing the ALife test: activity statistics classify evolution in Geb as unbounded.” *Advances in Artificial Life: Proceedings of the Sixth European Conference on Artificial Life (ECAL2001)*, Kelemen, J, and Sosik, P (eds), pp. 417–426, Springer, Heidelberg, Germany.
- Cope, ED (1871). “The method of creation of organic forms.” *Proc. Am. Philos. Soc.* **12**, 229–263.
- Dawkins, R (1997). “Human chauvinism (a review of Gould's Full House).” *Evolution* **51**(3), 1015–1020.
- Dennett, DC (1996). *Kinds of Minds: Toward an Understanding of Consciousness*, Basic Books, New York, NY.
- Farmer, JD, and Griffith, V (2007). *On the Viability of Self-Reproducing Machines* (unpublished).
- Fretwell, SD (1972). *Populations in a Seasonal Environment*, Princeton University Press, Princeton, NJ.
- Fretwell, SD, and Lucas, HJ (1970). “On territorial behavior and other factors influencing habitat distribution in birds. I. Theoretical development.” *Acta Biotheor.* **19**, 16–36.
- Gould, SJ (1989). *Wonderful Life: The Burgess Shale and the Nature of History*, Norton, New York.
- Gould, SJ (1994). “The evolution of life on earth.” *Sci. Am.* **271**(4), 62–69.
- Gould, SJ (1996). *Full House*, Harmony Books, New York.
- Griffith, V, and Yeager, LS (2006). “Ideal free distribution in agents with evolved neural architectures.” *Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, Rocha, L, Yeager, LS, Bedau, M, Floreano, D, Goldstone, R, and Vespignani, A (eds), pp. 372–378, MIT Press (Bradford Books), Cambridge, MA.
- Heylighen, F (2000). “Evolutionary transitions: how do levels of complexity emerge?” (citeseer.ist.psu.edu/375313.html)
- Huynen, MA (1996). “Exploring phenotype space through neutral evolution.” *J. Mol. Evol.* **43**(3), 165–169.
- Katz, MJ (1987). “Is evolution random?” *Development as an Evolutionary Process*, Raff, RA, and Raff, EC (eds), pp. 285–315, Alan R. Liss, New York.
- Kimura, M (1983). *The Neutral Theory of Molecular Evolution* Cambridge University Press, Cambridge, UK.
- Knoll, AH, and Bambach, RK (2000). Directionality in the history of life:

- diffusion from the left wall or repeated scaling of the right?" *Deep Time: Paleobiology's Perspective*, Erwin, DH and Wing, SL (eds), Vol. 26 (sp4), 1–14.
- Lizier, JT, Piraveenan, M, Dany, P, Prokopenko, M, and Yaeger, LS (2009). "Functional and structural topologies in evolved neural networks." *Advances in Artificial Life: Proceedings of the Tenth European Conference on Artificial Life (ECAL2009)*, Kampis, G, Szathmry, E, Fernando, C, Jelasity, M, Jordn, F, Lrincz, A, and Scheuring, I (eds), Springer, Heidelberg, Germany.
- Lungarella, M, Pegors, T, Bulwinkle, D, and Sporns, O (2005). "Methods for quantifying the information structure of sensory and motor data." *Neuroinformatics* 3(3), 243–262.
- Maynard Smith, J (1970). "Time in the evolutionary process." *Stud. Gen. (Berl.)* 23, 266–272.0039-4149
- McShea, DW (1993). "Evolutionary change in the morphological complexity of the mammalian vertebral column." *Evolution* 47(3), 730–740.
- McShea, DW (1994). "Mechanisms of large-scale evolutionary trends." *Evolution* 48, 1747–1763.
- McShea, DW (1996). "Metazoan complexity and evolution: is there a trend?" *Evolution* 50, 477–492.
- McShea, DW (2001). "The minor transitions in hierarchical evolution and the question of a directional bias." *J. Evol. Biol.* 14, 502–518.
- McShea, DW (2005). "The evolution of complexity without natural selection, a possible large-scale trend of the fourth kind." *Paleobiology*, 31 (sp2), 146–156.
- Newman, MEJ, and Engelhardt, R (1998). "Effects of neutral selection on the evolution of molecular species." *Proc. R. Soc. London, Ser. B* 265, 1333–1338.
- Raup, DM, and Gould, SJ (1974). "Stochastic simulation and evolution of morphology—towards a nomothetic paleontology." *Syst. Zool.* 23(3), 305–322.
- Raup, DM, Gould, SJ, Schopf, TJM, and Simberloff, DS (1973). "Stochastic models of phylogeny and the evolution of diversity." *J. Geol.* 81, 525–542.
- Rechsteiner, A, and Bedau, MA (1999). "A generic neutral model for quantitative comparison of genotypic evolutionary activity." *Advances in Artificial Life: Fifth European Conference on Artificial Life Proceedings*, pp. 109–118, Springer, Heidelberg, Germany.
- Rensch, B (1960a). *Evolution Above the Species Level*, Columbia University Press, New York.
- Rensch, B (1960b). "The laws of evolution." *The Evolution of Life*, Tax, S (ed), pp. 95–116, University of Chicago Press, Chicago, IL.
- Saunders, PT, and Ho, MW (1976). "On the increase in complexity in evolution." *J. Theor. Biol.* 63, 375–384.
- Seth, AK, and Edelman, GM (2004). "Environment and behavior influence the complexity of evolved neural networks." *Adapt. Behav.* 12(1), 5–20.
- Shannon, CE (1948). "A mathematical theory of communication." *Bell Syst. Tech. J.*, 27, 379–423, 623–656.
- Sporns, O, and Kötter, R (2004). "Motifs in brain networks." *PLoS Biol.* 2(11), e369.
- Sporns, O, and Lungarella, M (2006). "Evolving coordinated behavior by maximizing information structure." *Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, Rocha, L (ed), pp. 323–329, MIT Press (Bradford Books), Cambridge, MA.
- Sporns, O, Tononi, G, and Edelman, G (2000). "Theoretical neuroanatomy: relating anatomical and functional connectivity in graphs and cortical connection matrices." *Cereb. Cortex* 10, 127–141.
- Thomas, RDK, and Reif, W-E (1993). "The skeleton space: a finite set of organic designs." *Evolution* 47(2), 341–360.
- Tononi, G, Edelman, G, and Sporns, O (1998). "Complexity and coherency: integrating information in the brain." *Trends Cogn. Sci.* 2(12), 474–484.
- Tononi, G, Sporns, O, and Edelman, G (1994). "A measure for brain complexity: relating functional segregation and integration in the nervous system." *Proc. Natl. Acad. Sci. U.S.A.* 91, 5033–5037.
- Turney, P (1999). "Increasing evolvability considered as a large-scale trend in evolution." *Proceedings of the Workshop on Evolvability at the 1999 Genetic and Evolutionary Computation Conference (GECCO-99)*, Wu, A (ed), pp. 43–46.
- Turney, P (2000). "A simple model of unbounded evolutionary versatility as a largest-scale trend in organismal evolution." *Artif. Life* 6, 109–128.
- Valentine, JW, Collins, AG, and Meyer, CP (1994). "Morphological complexity increase in metazoans." *Paleobiology* 20(2), 131–142.
- Waddington, CH (1969). "Paradigm for an evolutionary process." In *Towards a Theoretical Biology*, Waddington, CH (ed), Vol. 2, pp. 106–128, Aldine, Chicago, IL.
- Wagner, PJ (1996). "Contrasting the underlying patterns of active trends in morphologic evolution." *Evolution* 50(3), 990–1007.
- Yaeger, LS (1994). "Computational genetics, physiology, metabolism, neural systems, learning, vision, and behavior or polyworld: life in a new context." *Proceedings of the Artificial Life III Conference*, Langton, CG (ed), pp. 263–298, Addison-Wesley, Reading, MA.
- Yaeger, LS, Griffith, V, and Sporns, O (2008). "Passive and driven trends in the evolution of complexity." *Artificial Life XI: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, Bullock, S, Noble, J, Watson, R, and Bedau, MA (eds), pp. 725–732, MIT Press, Cambridge, MA.
- Yaeger, LS, and Sporns, O (2006). "Evolution of neural structure and complexity in a computational ecology." *Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, Rocha, L, Yaeger, LS, Bedau, M, Floreano, D, Goldstone, R, and Vespignani, A (eds), pp. 330–336, MIT Press (Bradford Books), Cambridge, MA.