

Passive and Driven Trends in the Evolution of Complexity

Larry Yaeger¹, Virgil Griffith^{1,2}, and Olaf Sporns³

¹School of Informatics and ³Department of Psychological and Brain Sciences
Indiana University, Bloomington, IN 47408

²Computation & Neural Systems, California Institute of Technology, Pasadena, CA 91125
larry@indiana.edu

Abstract

The nature and source of evolutionary trends in complexity is difficult to assess from the fossil record, and the *driven vs. passive* nature of such trends has been debated for decades. There are also questions about how effectively artificial life software can evolve increasing levels of complexity. We extend our previous work demonstrating an evolutionary increase in an information theoretic measure of neural complexity in an artificial life system (Polyworld), and introduce a new technique for distinguishing driven from passive trends in complexity. Our experiments show that evolution can and does select for complexity increases in a driven fashion, in some circumstances, but under other conditions it can also select for complexity stability. It is suggested that the evolution of complexity is entirely driven—just not in a single direction—at the scale of species. This leaves open the question of evolutionary trends at larger scales.

Introduction

The existence of an evolutionary trend towards greater complexity is undeniable, whether one measures complexity by organism size (Cope, 1871), distinct cell types (Bonner, 1988; Valentine et al., 1994), morphology (Thomas and Reif, 1993; McShea, 1993), or ecological webs of interaction (Knoll and Bambach, 2000). Historically, it has often been suggested that such growth is the result of an evolutionary bias towards forms and functions of greater complexity, and a great variety of rationales has been offered for why this should be the case; e.g., (Rensch, 1960a,b; Waddington, 1969; Saunders and Ho, 1976; Kimura, 1983; Katz, 1987; Bonner, 1988; Arthur, 1994; Huynen, 1996; Newman and Engelhardt, 1998); see McShea (1991) and Carroll (2001) for reviews. However, Maynard Smith (1970), Raup et al. (1973), Gould (1994), and others have questioned whether that growth has been 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.

Bedau et al. (1997) and Rechsteiner and Bedau (1999) provide some evidence of an increasing and accelerating

"evolutionary activity" in biological systems not yet demonstrated in artificial life models. However, other attempts to characterize complexity trends in the fossil 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. McShea (1994, 1996, 2001, 2005) has, over the years, 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 asymmetric random drift.

Simple computational models of branching species and clade lineages in simple numerical parameter spaces (arbitrary values standing in for complexity, size, or the like) have been used to investigate this distinction between driven and passive evolutionary trends (Raup et al., 1973; Raup and Gould, 1974), however it is not always possible to distinguish a passive system from a weakly driven system (McShea, 1994). Furthermore, the anagenetic component of some of these models, while intended, by definition, to address within-lineage change, is equivalent to branching lineages that effectively compete with one another for parameter space, by requiring branched, descendant lines to replace ancestral lines. The typical assumption of equal extinction rates across all scales may also unintentionally color the results from these models. It is common wisdom, for example, that a driven system necessarily implies an increase in the minimum value of whatever parameter is being used to distinguish taxonomic branches (Wagner, 1996; McShea, 2001; Carroll, 2001), yet an evolutionary system in which fitness at smaller scales is independent of fitness at larger scales could possess a drive towards larger scales without eliminating or even disadvantaging organisms at the lower end of the spectrum. Indeed, McShea (1994) acknowledges a dramatically lower rate of growth in this minimum in a purely cladogenetic model compared to a mixed anagenetic and cladogenetic model. One of us [LY] is currently investigating the effects of the uniform extinction rate assumption.

Given the difficulty and ambiguity one encounters when

attempting to answer questions about the evolution of complexity from paleontological data or simple branching models, it makes sense to turn to computer models of evolution to address these questions. Turney (1999, 2000) has 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); Adami (2002) has defined complexity as the information that an organism's genome encodes about its environment and used Avida to show that asexual agents in a fixed, single niche always evolve towards greater complexity of this narrowly defined type.

Modern compute power and artificial life methods allow us to rewind the "tape of life", as Gould (1989) put it, and let history unfold again and again under slightly or dramatically different influences. Here we use a method for "replaying the tape" that is substantially beyond the perturbed playback Gould envisaged, demonstrating a method for carrying out parallel simulations in which natural selection either does or does not play a part, yet with all other population and genetical statistics being held constant. This is similar in spirit to the extreme "behavioral noise" null model used by Bedau (1995) and subsequent "neutral shadow models" (Bedau et al., 1998). Being able to effectively turn natural selection on and off in this fashion allows us to tease apart and distinguish evolutionarily driven trends from passive trends in a formal, quantitative fashion.

The trend we investigate is a particular information-theoretic measure of complexity (Tononi et al., 1994; Lungarella et al., 2005), C , for the neural dynamics of artificial agents in an evolving computational ecology, Polyworld (Yaeger, 1994). In previous work (Yaeger and Sporns, 2006) we demonstrated an increasing trend in C in the agents of Polyworld over evolutionary time scales, and were able to relate these increases to increasing structural elaboration of the agents' neural network architectures and an increase in the learning rates employed at the Hebbian synapses in these networks. We did not, however, address the evolutionary source of these increases, or whether that source should be construed as driven or passive, in the McShea (1996) sense. Here we use a novel technique that allows us to make such a distinction, and discover that, at least at the scale of single species and ecological niches, evolution of complexity is always driven, but, interestingly, not always driven in the same direction.

Tools and Techniques

Polyworld

Polyworld (Yaeger, 1994) is an evolutionary model of an ecology populated with haploid agents, each with a suite of primitive behaviors (move, turn, eat, mate, attack, light, focus) under continuous control of an Artificial Neural Network (ANN) consisting of summing and squashing neurons

and synapses that adapt via Hebbian learning. The architecture of the ANN is encoded in the organism's genome, expressed as a number of neural groups of excitatory and inhibitory neurons, with genetically determined synaptic connection densities, topologies, and learning rates. Input to the ANN consists of pixels from a rendering of the scene from each agent's point of view, like light falling on a retina. Though agent morphologies are simple and static, agents interact with the world and each other in fairly complex ways, as they replenish energy by seeking out and consuming food and by attacking, killing, and eating other agents. They reproduce by simultaneous expression of a mating behavior by two collocated agents.

Agent population is normally bounded above and below, but unlike the simulations discussed in (Yaeger and Sporns, 2006), there is no "smite" function invoked at maximum population, which the authors felt risked introducing a bias associated with the simulator's *ad hoc* heuristic fitness function. Nor does a minimum population any longer invoke a steady-state GA, which also would necessarily depend upon that fitness function. Instead, as the population grows towards the upper bound, the amount of energy depleted by all agent behaviors, including neural activity, is increased in a continuous fashion. Conversely, as the population drops towards the minimum, energy depletion is decreased, and agent lifespans may be artificially extended. This does not guarantee a viable population (one that sustains its numbers through reproduction), since unsuccessful, unfit agents may be all that remain by the time a failing population bottoms out, but it does provide very effective population control without denying births to agents capable and "desirous" of doing so, while simultaneously eliminating any possible effects of the now purely informative heuristic fitness function.

Complexity

For our purposes, complexity, C , is computed using a new C++ implementation based on the methods of (Lungarella et al., 2005) for approximating the information-theoretic measure of complexity originally developed by (Tononi et al., 1994). Though non-trivial to derive and implement, the intuition behind C is straightforward: Cooperation amongst 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* and measured by the difference between maximum and actual integration at different scales, also increases network complexity. Maximal complexity is achieved in networks that optimally trade off the opposing tensions between integration and segregation—between cooperation and specialization—and maximize both to the extent possible. The original mea-

sure of complexity is given by:

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

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 we use was introduced in (Tononi et al., 1998) and explored computationally in (Sporns et al., 2000):

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

where $H(X)$ is 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 entropy of the rest of the system.

Natural Selection vs. Random Drift

While probably impossible to eliminate natural selection from an evolving biological ecology, artificial ecologies are more flexible. In order to distinguish between evolutionarily driven and passive trends, we designed a mode for running our simulator in which natural selection had effectively been eliminated, yet which could be compared directly with runs in which natural selection operated normally. This was accomplished by implementing a new “lockstep” mode of operation in Polyworld. First a simulation is run in the system’s normal, natural-selection mode of operation. During this natural-selection run, the birth and death of every agent is recorded (along with the usual statistics, brain states, etc.). Then the simulator is run in the lockstep mode, starting from the same initial conditions as the natural-selection run and using the birth and death data recorded during the natural-selection run. No “natural” births or deaths are allowed during a lockstep run. Instead, every time a birth occurred in the original natural-selection run, a birth is forced to occur in the lockstep run, only instead of being produced by the original parents, the birth is produced by two agents chosen at random from the population. Similarly, whenever a death occurred in the natural-selection run, a death is forced to occur in the lockstep run, only instead of the original agent dying, a random agent is killed and removed from the population.

By so doing, population statistics are forced to be identical between the paired natural-selection and lockstep runs. As a result, the genetical statistics—number of crossover and mutation operations—are forced to be comparable in the paired runs. Note that since crossover and mutation are applied to different genomes and since the number of crossover points and the mutation rate are themselves embedded in these genomes (Yaeger, 1994), these genetic operations are

only statistically comparable between paired runs, not identical. 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 impact the values of neural complexity we compute, this could produce extraneous differences between paired runs, but we do not expect this to have any consistent, measurable influence. The controlling statistics, such as the entropy and mutual information in the visual inputs, are comparable between paired runs, so we expect computed complexity values to be similarly comparable. While it would be possible to record number of crossover points, mutation rates, agent trajectories, and even sensory inputs during the natural-selection runs and play them back during the lockstep runs, we do not believe this would alter the relevant statistics or the measured outcomes in a substantive manner, and therefore have not made any such attempts. One could also argue that since complexity is affected by agent behaviors and their resulting sensory inputs, agents in lockstep runs must be able to control their actions in order to obtain valid measures of their neural complexity.

The end result of these machinations is that gene states are subject to natural selection, based on the evolutionary viability—the fitness—of the agents’ behaviors, in the natural-selection runs. While gene states are subject only to the same degree of variation, with no evolutionary fitness consequences or effects, in the lockstep runs. Additionally, population statistics are identical and sensory input statistics are comparable between paired runs.

Simulations and Data Acquisition

A set of 10 paired simulations, differing only in initial random number seed, were run in natural-selection and lockstep modes; i.e., 20 simulations in all. Each was run for 30,000 time steps. As Polyworld is continuous rather than generational, determining the number of generations is non-trivial. In the past estimates have been known only to fall within a large range. A low estimate based on average lifespan (about 300 time steps) would be 100 generations. A high estimate based on the minimum age of fecundity (25 time steps) would be 1,200 generations. A newly implemented lineage tracer produces a more accurate estimate of about 400 generations.

The world is seeded with a uniform population of agents that have the minimum number of neural groups and nearly minimal neuron and synapse counts. While predisposed to some potentially useful behaviors, such as running towards green (food) and away from red (aggressive agent behaviors; see (Yaeger, 1994) for details on color use in Polyworld), these seed organisms are not a viable species. That is, unless they evolve they cannot sustain their numbers through reproduction and will gradually die out.

During simulations, the activation of every neuron in the

brain of every agent is recorded at every time step. These brain function recordings are grouped into arbitrary (here, 1,000) time step bins, for all agents that died during the specified interval. Utility programs are then used to calculate the complexity, C , of the neural dynamics of every agent's complete lifespan (hence the requirement for the agent's death). We then compute mean complexities for these binned populations of agents as a function of time. Finally, we compute means and standard deviations of the population means for the multiple natural-selection and lockstep simulations as a function of time to study general evolutionary trends in complexity.

Complexity can be calculated across all neurons, just the input neurons, or just the "processing" neurons (all neurons except inputs). All complexities presented here are based on processing neurons. (In general, there is little difference in complexity trends between all neurons and processing neurons.) Complexity varies as evolution produces changes in the parts of the genome that specify the neural architecture.

Results and Discussion

Figure 1 shows complexity versus time for the previously described series of 10 paired driven (natural-selection) and passive (lockstep) simulations. The lighter lines depict population means from individual runs. The heavier lines depict means of all runs of a particular type (driven or passive). Data is presented in this individual-plus-mean fashion, rather than mean-plus-standard-error fashion, to give a better feel for the nature of the variance between runs, and to identify some interesting events in a small number of the runs (discussed later). Plotted beneath the complexity lines is a single dotted line that measures a paired or dependent Student's T-test computed on the same time interval as the complexity data are computed and plotted. (A dependent test is used because these are paired runs with common initial conditions and enforced common population and genetic statistics.) Where this line is above the horizontal T-critical (T^*) line, this standard measure of statistical significance rejects the null hypothesis with $p < 0.05$; at or below the T-critical line the null hypothesis cannot be rejected (at least not as reliably). Given 10 pairs of runs, the number of degrees of freedom is 9, and T-critical is 1.833.

The first thing to note is a statistically significant faster growth rate in complexity in the driven runs than in the passive runs during approximately the first 4,000 time steps. Evolution is clearly selecting for an increase in complexity during this early time period. This makes sense intuitively since the seed population is known to be non-viable and must evolve or die out. Increases in complexity during this period are of a distinct evolutionary advantage, producing descendant populations that are more capable of thriving in this particular environment than their ancestors. During this time period, the evolution of complexity is clearly driven, with a bias towards increasing complexity.

The next thing to note is the early plateauing of complexity in the driven runs, allowing the randomly drifting complexity of the passive runs to catch and surpass them by around $t=7,000$. This is the result of evolution having found a solution that is "good enough", and the concomitant spread of the genes producing this solution throughout the population. Seven out of the 10 natural selection simulations remain relatively stable around this modestly complex solution once it is found. The intuition here is that any change away from this "good enough" solution is likely to be detrimental, hence evolution selects for stability. Note that this actively suppresses genetic drift and, indeed, a statistically significant difference between driven and passive runs, with passive complexity now being the larger of the two, is maintained from about $t=8,000$ to the end of the runs at $t=30,000$.

There is also a consistent, but less interesting, plateauing of complexity in the passive runs. This is due solely to the individual bits of the underlying genome approaching a state of approximately 50% on, 50% off. Effectively, the random walk has maximized variance as much as it can given the model parameters. Though generally higher than the driven mean, complexities in the passive/random model are nowhere near the maximum obtainable with the full range of gene values (as observed in the complexity-as-fitness-function experiments discussed below); they just correspond to the range of complexities representable by the genome with an even mix of on and off bits. Such larger values of complexity are potentially meaningful, but do not confer any evolutionary advantage on agents in these lockstep runs.

Finally, if one looks carefully at driven complexity for the individual runs, three (of the 10) runs make secondary transitions upward in complexity between $t=20,000$ and the end of the run, coincidentally reaching about the same level of complexity as the passive runs. In this subset of runs, apparently a new or improved behavior emerges late in the simulation, and the genes producing this behavior spread throughout the population fairly rapidly. Despite the simplicity of the world design used for these experiments, multiple viable, competing solutions have emerged and it is always possible that more solutions would emerge given time. Importantly, for the future, it appears that this complexity measure provides a useful tool for ascertaining the onset of new, improved strategies, including speciation events, as well as a quantitative tool for assessing the neural changes that produced the new strategies.

Two of us [LY,OS] independently realized some years ago that if we were in a position to measure neural complexity in an artificial system, it might make sense to accelerate the evolution of complexity by using it as an explicit fitness function. Although not elaborated upon here, preliminary experiments to this end have been carried out. With the same values for the parameters controlling neural architecture, using complexity as a fitness function pushes mean complexity up to around 0.9, roughly three times the levels




Figure 1: Driven and passive complexity vs. time. Light solid lines show population mean complexity for each driven, natural-selection run. Light dashed lines show population mean complexity for each passive, lockstep run. Heavy lines show means of all ten runs for corresponding line style. Light dotted line at bottom shows dependent Student's T-test relative to horizontal T-critical line (labeled T*) for $p > 0.05$.

obtained by this series of driven (0.275) and passive (0.33) runs. However, selecting purely for complexity consistently produces stereotypic spinning behaviors by the agents, that would not be of much value under natural selection, and we suspect are the result of a maximization of entropy and mutual information in the sensory inputs. This differs from the direct coupling between complexity and behavior found in (Sporns and Lungarella, 2006), probably due to differences in the range of possible behaviors and the nature of the sensory inputs in the simulation environments used for these two studies.

Adami et al. (2000); Adami (2002) has defined a measure of genomic complexity in terms of how much information a species' genome encodes about its environment, based on the cross-population entropy at each genomic site (here, individual bits). The measure has acknowledged constraints—it only applies to a single species in a single, static niche, thus failing to capture issues related to biodiversity, environmental variability, or broader ecologies. More generally, we suspect it may be a better measure of genomic consistency or specialization of a species than of complexity, but there is no question that the aggregate stability of gene sites in a species' population is an important measure of the success of that species at encoding information about its environment

in its genome. Since our current series of simulations are deliberately simple and may probably best be thought of as the evolution of a single (or at least highly related) species in a single, static niche, we decided to investigate the evolution of this genomic consistency.

Figure 2 shows the evolution of Genomic Consistency (GC) over time. Since the world is seeded with a uniform population, GC is initially extremely large, as the measure effectively assumes the current genetic structure is an evolved response to the environment and perfect uniformity is maximally consistent. (However, we do not feel this should be seen as maximally complex, hence our renaming of Adami's genomic complexity.) Accordingly, the vertical extent of the graph has been truncated in order to focus on the more interesting results. The main observation is the dramatic difference between driven and passive runs. The passive runs produce extensive random gene edits, thus minimizing GC. The driven runs demonstrate a larger, stabilized GC across the population over time due to natural selection for specific traits that increase the evolutionary fitness of the agents. There is a hint of a modest upward trend in GC after it bottoms out around $t=12,000$, suggesting a possible continued incorporation of information about the environment in the genome of these agents, but so far no attempt has been




Figure 2: Genomic consistency vs. time. Light solid lines are for each driven, natural-selection run. Light dashed lines are for each passive, lockstep run. Heavy lines show means of all ten runs for corresponding line style.

made to establish statistical significance.

Another way to draw some understanding from this series of driven vs. passive runs is to look at the distribution of complexity in the populations over time. Figure 3 shows time-series histograms of complexity throughout the population for four sample simulations—two passive, two driven. The two passive runs show a generalized increase in variance, due to the diffusive random walk away from the low complexity of the seed agents. The two driven runs show a more peaked distribution around the complexity values attained by the viable populations emerging as a result of natural selection. Figure 3(c) is representative of the majority of the driven runs, showing a shift towards a modest level of complexity. Figure 3(d) is representative of the small number of driven runs in which a secondary transition to a different behavior and higher level of complexity emerged late in the simulation.

Conclusions

We have demonstrated a technique for directly comparing and assessing neural complexity growth in equivalent driven and passive systems. Using this technique we have shown evolutionary selection for *increased* complexity, in a “driven” fashion, as well as selection for complexity *stability*. Though we have not demonstrated it here, there is little doubt that a system in which the cost of neural com-

plexity exceeded its value would result in a driven *reduction* in complexity, the way dark dwelling organisms in a cave have been known to give up their eyes. This paints a complex picture of evolutionary selection for increasing, stable, and decreasing complexity, none of which corresponds to a purely “passive” mechanism of complexity change. At this scale, evolution is entirely driven, with changes in complexity always being selected for or against. Scale, however, is very important to this discussion.

Gould (1996) and Dawkins (1997) have argued strongly for passive and driven evolutionary trends, respectively. However, much of the disconnect between them seems to be precisely an issue of scale. Dawkins is unquestionably correct about evolution being driven on a short time scale, for a small set of related species. Yet Gould *may* be correct, as well, about evolution being fundamentally passive on a longer time scale, over the entire tree of biological life. In one of the earliest works to model evolution computationally in order to characterize active versus passive trends, Raup et al. (1973) called attention to the fact that fully deterministic, driven trends acting at small scales are in fact likely to be at the base of larger scale trends, even if those large scale trends turn out to be passive. A mix of many, potentially opposing trends might very well appear random and undirected when integrated together. What our current simulations show is that, indeed, while evolution undoubtedly




Figure 3: Histograms of complexity over time for individual runs. (a) and (b) are passive runs. (c) and (d) are driven runs.

drives complexity changes, according to perfectly standard expectations about the evolutionary fitness of those changes, it does not drive in just one direction. When complexity increase is of an evolutionary advantage it will be selected for, just as will complexity decrease. And when a species' complexity is "good enough", so that any increase or decrease is likely to involve a step away from a local fitness maximum, evolution will mildly select for and stabilize the ex-

isting level of complexity. This goes a long way towards explaining the observation by Dennett (1996), "The cheapest, least intensively designed system will be 'discovered' first by Mother Nature, and myopically selected."

Looking forward, though we have yet to address the issue experimentally, we expect any increase in agent interactions with the world, any increase in complexity of the environment, and any increase in the available range of niches—Knoll and Bambach (2000)'s expanding *ecospace*—to produce an increase in evolved neural complexity of agents in the world. All niches are not created equal, and we suspect that evolutionary occupation of more and richer parts of *ecospace* will, as Knoll suggests, result in a fundamentally driven growth in complexity both at the largest scales of biology and in our artificial worlds. And in Polyworld we expect **C**, our measure of neural complexity, to quantify and document that trend.

Acknowledgements

Thanks to David Brinda for Python and gnuplot hacking.

References

- Adami, C. (2002). What is complexity? *BioEssays*, 24:1085–1094.
- Adami, C., Ofria, C., and Collier, T. (2000). Evolution of biological complexity. *PNAS*, 97(9):4463–4468.
- Arthur, B. (1994). On the evolution of complexity. In Cowan, G. A. e. a., editor, *Complexity: Metaphors, Models, and Reality*, pages 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. In Husbands, P. and Harvey, I., editors, *Proceedings of the Fourth European Conference on Artificial Life*, pages 125–134. MIT Press, Cambridge, MA.
- Bedau, M. A. (1995). Three illustrations of artificial life's working hypothesis. In Banzhaf, W. and Eeckman, F. E., editors, *Evolution and Biocomputation: Computational Models of Evolution, Lecture Notes in Computer Science Vol. 899*, pages 53–68. Springer, Berlin.
- Bedau, M. A., Snyder, E., and Packard, N. H. (1998). A classification of long-term evolutionary dynamics. In Adami, C. e. a., editor, *Artificial Life VI, Proceedings of the Sixth International Conference on the Simulation and Synthesis of Living Systems*, pages 228–237. MIT Press, Cambridge, MA.
- Bonner, J. (1988). *The Evolution of Complexity by Means of Natural Selection*. Princeton Univ. Press., Princeton, NJ.
- Carroll, S. (2001). Chance and necessity: the evolution of morphological complexity and diversity. *Nature*, 409:1102–1109.
- Cope, E. D. (1871). The method of creation of organic forms. *Proc. Am. Phil. Soc.*, 12:229–263.
- Dawkins, R. (1997). Human chauvinism (a review of gould's full house). *Evolution*, 51(3):1015–1020.
- Dennett, D. C. (1996). *Kinds of Minds: Toward an Understanding of Consciousness*. HarperCollins.

- Gould, S. J. (1989). *Wonderful Life: The Burgess Shale and the Nature of History*. Norton, New York.
- Gould, S. J. (1994). The evolution of life on earth. *Scientific American*, 271(4):62–69.
- Gould, S. J. (1996). *Full House*. Harmony Books, New York.
- Heylighen, F. (2000). Evolutionary transitions: how do levels of complexity emerge? citeseer.ist.psu.edu/375313.html.
- Huynen, M. A. (1996). Exploring phenotype space through neutral evolution. *Journal of Molecular Evolution*, 43(3):165–169.
- Katz, M. J. (1987). *Is evolution random?* Alan R. Liss, New York.
- Kimura, M. (1983). *The Neutral Theory of Molecular Evolution*. Cambridge University Press.
- Knoll, A. and Bambach, R. K. (2000). Directionality in the history of life: Diffusion from the left wall or repeated scaling of the right? *Paleobiology*, 26:4 (Supplement):1–14.
- 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. *Studium Generale*, 23:266–272.
- McShea, D. W. (1991). Complexity and evolution: What everybody knows. *Biology and Philosophy*, 6:303–324.
- McShea, D. W. (1993). Evolutionary change in the morphological complexity of the mammalian vertebral column. *Evolution*, 47(3):730–740.
- McShea, D. W. (1994). Mechanisms of large-scale evolutionary trends. *Evolution*, 48:1747–1763.
- McShea, D. W. (1996). Metazoan complexity and evolution: Is there a trend? *Evolution*, 50:477–492.
- McShea, D. W. (2001). The minor transitions in hierarchical evolution and the question of a directional bias. *J. Evol. Biol.*, 14:502–518.
- McShea, D. W. (2005). The evolution of complexity without natural selection, a possible large-scale trend of the fourth kind. *Paleobiology*, 31:2 (Supplement):146–156.
- Newman, M. E. J. and Engelhardt, R. (1998). Effects of neutral selection on the evolution of molecular species. *Proc. R. Soc. London B*, pages 1333–1338.
- Raup, D. M. and Gould, S. J. (1974). Stochastic simulation and evolution of morphology—towards a nomothetic paleontology. *Systematic Zoology*, 23(3):305–322.
- Raup, D. M., Gould, S. J., Schopf, T. J. M., and Simberloff, D. S. (1973). Stochastic models of phylogeny and the evolution of diversity. *Journal of Geology*, 81:525–542.
- Rechsteiner, A. and Bedau, M. A. (1999). A generic neutral model for quantitative comparison of genotypic evolutionary activity. In *Advances in Artificial Life: Fifth European Conference on Artificial Life Proceedings*, pages 109–118. Springer-Verlag, Heidelberg.
- Rensch, B. (1960a). *Evolution above the species level*. Columbia Univ. Press, New York.
- Rensch, B. (1960b). The laws of evolution. In Tax, S., editor, *The Evolution of Life*, pages 95–116. Univ. of Chicago Press, Chicago.
- Saunders, P. T. and Ho, M. W. (1976). On the increase in complexity in evolution. *J. Theor. Biol.*, 63:375–384.
- Sporns, O. and Lungarella, M. (2006). Evolving coordinated behavior by maximizing information structure. In Rocha, L. e. a., editor, *Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, pages 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. *Cerebral Cortex*, 10:127–141.
- Thomas, R. D. K. 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 in Cognitive Sciences*, 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. Nat. Acad. Sci.*, 91:5033–5037.
- Turney, P. (1999). Increasing evolvability considered as a large-scale trend in evolution. In Wu, A., editor, *Proceedings of the Workshop on Evolvability at the 1999 Genetic and Evolutionary Computation Conference (GECCO-99)*, pages 43–46.
- Turney, P. (2000). A simple model of unbounded evolutionary versatility as a largest-scale trend in organismal evolution. *Artificial Life*, 6:109–128.
- Valentine, J. W., Collins, A. G., and Meyer, C. P. (1994). Morphological complexity increase in metazoans. *Paleobiology*, 20(2):131–142.
- Waddington, C. H. (1969). Paradigm for an evolutionary process. In Waddington, C. H., editor, *Towards a Theoretical Biology, Vol. 2*, pages 106–128. Aldine, Chicago.
- Wagner, P. J. (1996). Contrasting the underlying patterns of active trends in morphologic evolution. *Evolution*, 50(3):990–1007.
- Yaeger, L. S. (1994). Computational genetics, physiology, metabolism, neural systems, learning, vision, and behavior or Polyworld: Life in a new context. In Langton, C. G., editor, *Proceedings of the Artificial Life III Conference*, pages 263–298. Addison-Wesley.
- Yaeger, L. S. and Sporns, O. (2006). Evolution of neural structure and complexity in a computational ecology. In Rocha, L. e. a., editor, *Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, pages 330–336. MIT Press (Bradford Books), Cambridge, MA.