Evolution of Neural Structure and Complexity in a Computational Ecology

Larry S. Yaeger[1], Olaf Sporns[2]

[1]School of Informatics and [2]Department of Psychological and Brain Sciences
Indiana University, Bloomington, Indiana 47405
larryy@indiana.edu

Abstract

We analyze evolutionary trends in artificial neural dynamics and network architectures specified by haploid genomes in the Polyworld computational ecology. We discover consistent trends in neural connection densities, synaptic weights and learning rates, entropy, mutual information, and an information-theoretic measure of complexity. In particular, we observe a consistent trend towards greater structural elaboration and adaptability, with a concomitant and statistically significant growth in neural complexity.

Introduction

The existence of evolutionary trends in complexity is a much debated topic. Some (Gould 1994) have argued against any such trend. Others (Bonner 1988, McShea 1996, Heylighen 1999, Adami et al. 2000, Carroll 2001, Adami 2002) have presented evidence and interpretations that support a trend towards increasing complexity in biological and genomic evolution. Darwin (1871) argued for a gradualist interpretation of the emergence and evolution of intelligence that suggests an increasing arrow of complexity. Bedau (et al. 1997, Rechsteiner and Bedau 1999) has argued for evidence of an increasing and accelerating "evolutionary activity" in natural biological systems not demonstrated in artificial life models to date.

Previously, Langton (1990) showed that average mutual information (AMI) provides an effective measure of complexity in Cellular Automata (CA).[1] In particular, AMI peaks where complexity is greatest, correlating with Langton's lambda critical and Wolfram's (1984) Class IV, "complex" CA.[2] Crutchfield and Young (1990), amongst other more general, but less computable, formulations, define a statistical complexity for discrete finite automata (DFA), based on "symmetries" (any statistical regularities) induced by the DFA, that captures such a peak in complexity for DFA predicting binary strings. They relate their complexity measure to AMI between past and future sequences, but note that AMI is at most an upper bound on their DFA complexity. Shalizi (2003) has derived a related, pointwise "local statistical complexity", applicable to networks, that is based on mutual information between the "causal state" at a point in a graph and the past "light-cone" at that point.

Tononi et al. (1994) proposed an information-theoretic measure of complexity in networks that combines a measure of global integration (reduced independence, related to AMI) with a measure of segregation (specialization) of form and function. The emergence of such complexity and of specific structural features of networks is interrelated (Sporns et al., 2000). Complex neural dynamics requires specific structural motifs, as well as the active structuring of sensory inputs by embodied agents (Lungarella et al. 2005).

To study the general issue of evolutionary trends in complexity, and the specific issue of quantifying complexity in evolved ANNs, we have updated and extended an early artificial life model, Polyworld (Yaeger 1994), to permit the acquisition of anatomical and functional data from each agent's artificial neural network (ANN), from birth to death. We analyze evolutionary trends in the structure and function of these networks, using both connectivity and information-theoretical measures. By so doing, we seek to establish baseline quantitative measures of complexity that can be used to assess evolutionary trends in natural and artificial living systems.

Tools and Techniques

Polyworld Simulator

Polyworld is a computational 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 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. These network designs are expressed in a probabilistic, generative manner, so a single genome can produce a small class of neural

[1] Mitchell et al. (1993) refute some claims by Packard (1988) relating Langton's lambda parameter and computational capacities to evolved CA dynamics, but their only substantive observation about Langton's results was that for finite automata the variance in complexity, as a function of lambda is large (particularly near lambda critical), thus limiting the value of lambda as a predictor of the dynamics of any particular CA, even while it effectively predicts average CA dynamics.

[2] Gray (2003) argues that Wolfram's four classes may need to be augmented to fully specify the range of CA dynamics, but the original four classes are widely acknowledged to make a useful, if coarse, first cut.

anatomies, while the full range of genetic encodings can produce a very large class of neural anatomies, which are the main focus of evolution in Polyworld. Thus Polyworld offers the opportunity to explore the relationship between adaptive behavior, structural elaboration of networks, and agents' capacity to generate and integrate information, while offering complete access to all neural and ecological variables.

The primary input to the agents' ANNs, and sole sense mechanism, is vision. The 3D environment in which the simulation takes place is rendered (with traditional computer graphics imaging techniques) from each agent's point of view and the resulting pixel map is fed as input to the ANN, as if it were light falling on a retina.

To be successful and thrive in Polyworld agents must replenish energy expended on all activity by finding and eating food, or by killing and eating other agents. They must also find mating partners and reproduce. Since reproduction is under behavioral control of the agents, populations may be self-sustaining or not. When they *are* self-sustaining, there is no fitness function save natural selection. When a population is *not* self-sustaining, a steady-state Genetic Algorithm (GA) guarantees a minimum number of agents in the world, using an *ad hoc* heuristic fitness function that rewards reproduction, long life, consumption of food, and movement, to varying degrees. Fitness only affects selection and reproduction when the agents are not self-sustaining. Half of this series of simulations never invoked the fitness function for agent creation; of the remaining, the latest time the fitness function was used was t=3,263 (out of 10,000).

To limit computational demands, a maximum population of 300 agents was imposed in all simulations. Once this limit is reached, reproduction is halted until a death occurs, unless the "smite" option is enabled (see below). This limit can have profound consequences, eliminating or slowing evolutionary change due to the inability of viable agents to produce offspring. In these simulations, max population was reached between t=2,000 and t=5,000. The effects of this constraint (and of the smite functionality, see below) are sufficiently profound that all data reported herein indicate (by a gray bullet) the time at which maximum population was achieved.

To address this evolutionary slowdown, a "smite" option was added to Polyworld. When employed, as in these simulations, a list of least fit agents is maintained. Then when two agents would normally reproduce, were it not for the population cap, the currently least fit agent is culled, and the mating pair is allowed to reproduce, thus implementing a *replace least fit* GA technique. This approach appears to at least partially overcome the evolutionary slowdown resulting from the population limit.

Once engaged, however, this technique does affect the observed values of fitness. As a result of culling, agents have fewer opportunities to move, eat, and mate, and they die at a smaller fraction of their possible lifespan—all key contributors to heuristic fitness. Fitness is thus significantly diminished once smite is engaged. (Selection is not directly affected, since the steady-state GA is no longer in use by the time smite is engaged and natural selection is not based on heuristic fitness.) Max population marks on the graphs also effectively denote the time at which smite is engaged.

Details of the simulation engine can be found in (Yaeger 1994), or by consulting the open source code at http://sourceforge.net/projects/polyworld.

Network Analysis Methods

We apply two sets of measures—network and information theoretic—in analyzing the structure and function of evolving ANNs.

Simple network measures such as binary connection density, average connection weight, average activation per unit and total amount of synaptic change are gathered from connection matrices and activity traces of the 10 currently most fit individuals that are recorded at fixed time intervals of 1,000 time steps. Changes over the course of the lifetimes of these agents are collated by comparing connection patterns at birth and death (typically around 500 time steps apart).

Information theoretical measures are computed from complete records of neural activations over the lifetime of an agent, using computational methods and algorithms introduced in Lungarella et al. (2005). Entropy H(X) and mutual information MI(X,Y) are derived from discretized states x_i and y_i of the neural activation variables X and Y (traces of activation for two neurons), following standard formulae of statistical information theory:

$$H(X) = -\sum_{i=1}^{N_X} p(x_i)\log_2 p(x_i)$$

$$MI(X,Y) = H(X) + H(Y) - H(X,Y)$$

$$= -\sum_{i=1}^{N_X}\sum_{j=1}^{N_Y} p(x_i,y_j)\log_2 \frac{p(x_i)p(y_j)}{p(x_i,y_j)}$$

Entropy is calculated based on estimates for the probabilities of the N_X states of variable X, while mutual information results from distributions of single and joint state probabilities of variables X and Y, existing in N_X and N_Y discrete states, respectively. Time series were discretized using 16 states to ensure that single and joint probability histograms were sufficiently populated.

In addition to univariate and bivariate informational measures we also calculated measures that capture the distribution of information across an entire system of variables **X** (bold **X** denotes a set of N variables X_i; for our current purposes each X_i corresponds to one of the neural activation variables). We derived entropies for larger sets of variables from their covariance matrices using standard formulae (Papoulis 1990; Cover and Thomas 1991).

While MI(X,Y) quantifies statistical dependence of two variables, integration I(**X**) (Tononi et al. 1994) represents its multivariate extension measuring the total amount of statistical dependence between a set of variables **X**:

$$I(\mathbf{X}) = -\sum_{i=1}^{N} H(X_i) - H(\mathbf{X})$$

Complexity is incompatible with total integration (all variables maximally correlated or integrated) and with total randomness (all variables maximally decorrelated or segregated). Instead, complexity corresponds to the coexistence of integration and segregation within a single system (Tononi et al. 1994, Sporns et al. 2000). In neural terms, integration and segregation express the total amount of statistical dependence between neurons and the pattern of information distribution among them. A corresponding quantitative measure of complexity

$$C(\mathbf{X}) = H(\mathbf{X}) - \sum_{i=1}^{N} H(X_i \mid \mathbf{X} - X_i)$$

captures the degree to which globally integrated and locally segregated information coexists within the neural system **X**, composed of individual neurons X_i. ($\mathbf{X} - X_i$ is the set of all N-1 variables X_j where $j \neq i$.) Previous analyses have shown that C(**X**) is high for neural systems that effectively combine functional segregation and integration, e.g. by combining specialized neurons and global (system-wide) interactions within a common architecture. On the other hand, C(**X**) is low for random systems (lacking global integration), or for systems that are highly uniform (lacking local specialization).

Connectivity measures allow insights about the number, density, activation, and plasticity of neural units and their interconnections. Information-theoretical measures allow insights about the amount and distribution of information in input units (I) driven by sensory stimulation as well as interconnected processing units (P) that respond to inputs and generate motor outputs. All observed evolutionary changes are the result of progressive mutation, selection and variation in the agent population's genomes, as well as emergent processes related to behavioral interactions, population density, and the like.

Simulations and Data Acquisition

The Simulations

All of the simulations carried out for this analysis were identical except for the seed to the pseudo-random number generator. In all cases the world is configured so that food only grows in two bands of unequal area, but initially equal density, covering a total of 50% of the world area. This particular configuration is one sample from a series of runs investigating optimal foraging/resource matching (Griffith and Yaeger 2006), and is very simple in terms of what is required of the agents in order to survive and reproduce.

The world is initially seeded with 90 agents with identical seed genomes. (Their neural architectures vary subtly, due to the probabilistic nature of the expression of the genome when producing these neural architectures.) The seed genome was designed to elicit nominal, "reasonable" initial behaviors within the context of the Polyworld environment—move towards green (food), turn away from and attack red (attackers, barriers), mate with blue (agents expressing their mating behavior). The seed genome was crafted so as to produce a nearly viable species that exhibits easily understood behaviors. However, the initial seed population, without the benefit of variation and selection, cannot "succeed"—cannot sustain its numbers over time—with its existing range of neural anatomies and behaviors.

A seed genome was used in these experiments to confine evolutionary search to a smaller portion of the possible state space. (Similar runs using completely random initial seed populations produced viable populations, but the time to do so varied hugely, as did the range of solutions, making data gathering and data comparisons difficult.) To maintain this state space confinement, the steady-state GA creation of agents in low-population conditions was limited to producing matings between agents from the N-best list and to the occasional elitist reintroduction of the best ever genome. No random "hopeful monsters" were allowed.

Over slightly varying timescales, all ten simulation runs produced viable populations that matched the available resources in an approximate ideal free distribution (see Griffith and Yaeger 2006 for additional details).

The Data

As 10 simulations were carried out, the simulator recorded the following, for all agents:

- ANN anatomy—a signed, weighted connection matrix—at birth
- ANN anatomy at death
- ANN function—the neural activation levels for all neurons at all time steps of an agent's life

Data was culled on the fly, retaining only the top 10 "best" (based on heuristic fitness) agents that had died at any time in the previous 1,000 time steps. All runs lasted for 10,000 time steps. For analysis, neurons are divided into input (I) neurons, that relay sensory activation, and processing (P) neurons, including both internal and motor neurons.

In addition, a log of the mean and standard deviation, over the complete population, of the values of all genes, at each time step, was recorded during these runs. Subsequent data processing allowed the estimation of the average learning rate at all synapses in the ANNs over time.

Statistical trends in population data were evaluated using standard independent-measures t-tests for comparisons between the seed generation and the time point at which maximal population was reached (grey asterisk at t=3,500 indicates $p<0.01$), as well as between the seed generation and t=10,000 (black asterisk, $p<0.01$).

Results

Simulation Metrics

The simulation metric results are displayed in Figure 1.

Population. Populations reached the imposed maximum of 300 agents at times ranging from 2,000 to 5,000 time steps, averaging approximately 3,500 time steps (Fig. 1A). All 10 runs reached maximum population by 5,000 time steps. Reaching maximum population signifies an unequivocal success at surviving and thriving in the world. It also denotes the time at which the smite function is engaged.

Smitten. The number of agents smitten (see above) grows rapidly once the population reaches its maximum (Fig. 1B). This demonstrates how prolific the agents have become. At this time agent lifespan begins to decrease sharply, producing a downward trend in heuristic fitness.

Learning Rate. There is a pronounced growth in learning rate (Fig. 1C), as determined from the genes that control the Hebbian synaptic learning rate, averaged over the entire population at each time step. The growth is nearly linear over the life of the runs, unaffected by the population limit, smite activity, or heuristic fitness function, unlike some other metrics.

Fitness. On average, the heuristic fitness function averaged over the entire population at each time step (Fig. 1D) shows a strong increase from t=0 to t=3,000, followed by a plateau from t=3,000 to t=5,000, then a sharp decline from t=5,000 to simulation end. These transitions correspond to the average and final times at which populations reach the artificially imposed maximum, and thus the times at which the smite function is initially engaged and then comes to dominate population dynamics. The resulting reduction in lifespan produces a sharp drop in fitness. Early on, however, fitness continues to climb well after it has ceased to influence selection and reproduction (after about t=1,000), suggesting that however *ad hoc*, this heuristic function is capturing fitness in a way that is, at least initially, consistent with natural selection in this environment.

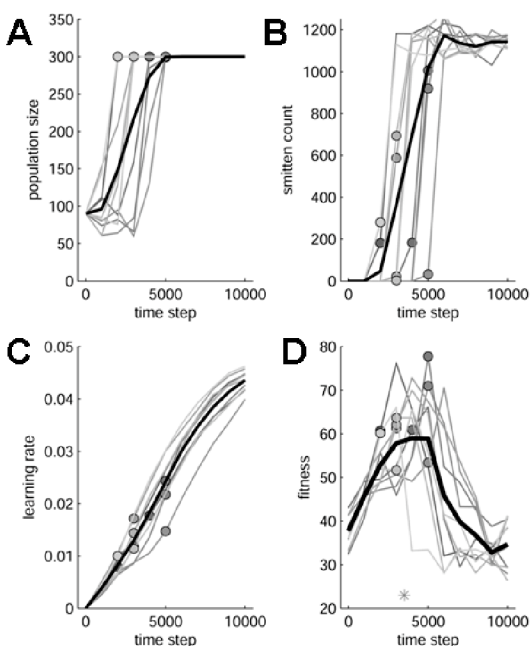

Figure 1: Simulation metrics. (A) Population count (B) Smitten count (C) Learning rate (D) Fitness.

Simple Network Metrics

Statistical trends of network metrics across agent populations are displayed in Figures 2 and 3.

Unit Number and Connection Density. While ANN sizes remained approximately constant for input and processing units (Fig. 2A,B), connection density underwent significant growth, with a faster rate among processing units (Fig. 2C,D). Average connection strengths increased between I and P neurons as well as amongst P neurons, the latter

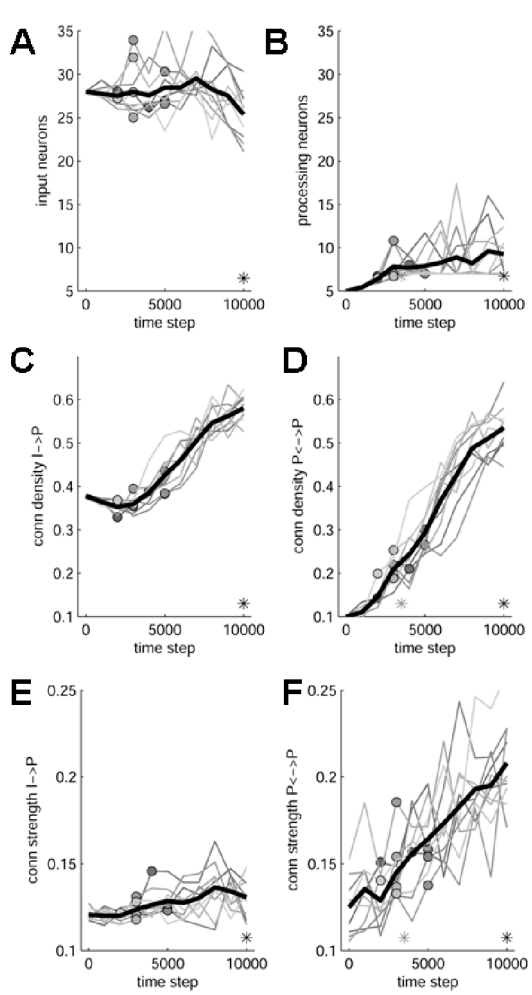

Figure 2: Network metrics. (A) Number of input (I) units (B) Number of processing (P) units (C) I→P connection density (D) P↔P connection density (E) Average absolute connection strengths for I→P connections and (F) P↔P connections.

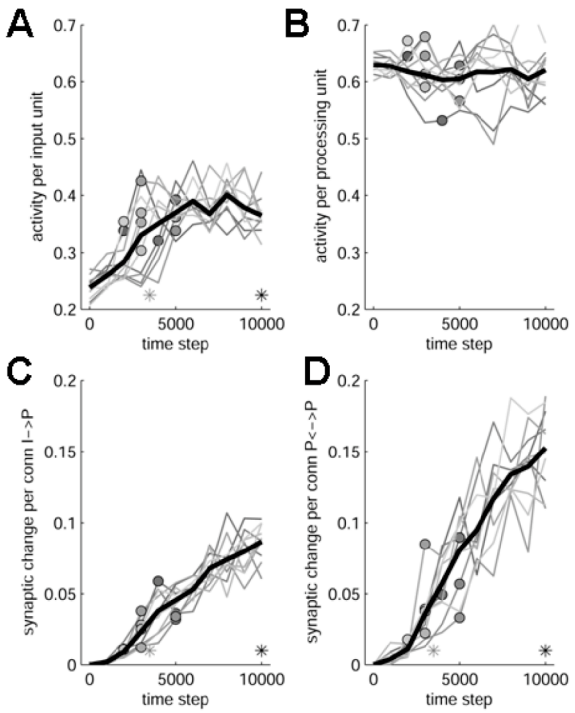

Figure 3: Network metrics (cont'd). (A) Average activity per unit per time step for I units and (B) P units. (C) Amount of synaptic change per connection for I→P and (D) P↔P connections.

showing earlier and larger increases overall (Fig. 2E,F), with much of this increase (measured at "death") due to learning (see below). Taken together, these changes indicate strong evolutionary pressure towards more structural connectivity within the agents' ANN.

Unit Activity. Despite these large increases in structural connectivity, changes in overall activation levels (expressed per neuron and per time step) over the course of evolution are relatively modest. I neuron activity initially rises somewhat (Fig. 3A), probably reflecting increases in population density and thus elevated levels of agent interactions. P neurons maintain constant activity levels (Fig. 3B), due to a preserved dynamical balance between excitatory and inhibitory connections (data not shown).

Synaptic Change. One of the most significant and unequivocal trends we observed was the persistent rise in lifetime learning over evolutionary time scales. Connections linking I to P neurons (Fig. 3C), as well as those linking P neurons among themselves (Fig. 3D), underwent increasing amounts of synaptic change (measured as the absolute difference between connection strengths at death and birth). This trend continues well after the smite function becomes fully engaged and would probably extend past t=10,000. This is undoubtedly due, in large part, to the evolved increase in learning rate, but the difference in slopes between I→P connections and P↔P connections (also reflected in Fig. 2E,F) suggests this is not the only mechanism at work and is not yet fully understood.

Information Theoretic Network Metrics

Information-theoretic measures are expressed on a per neuron basis (entropy), as averages over binary relationships between neurons (mutual information), or as global indices of the dynamics of an entire ANN (integration and complexity). Our results are shown in Figure 4 and discussed in detail below.

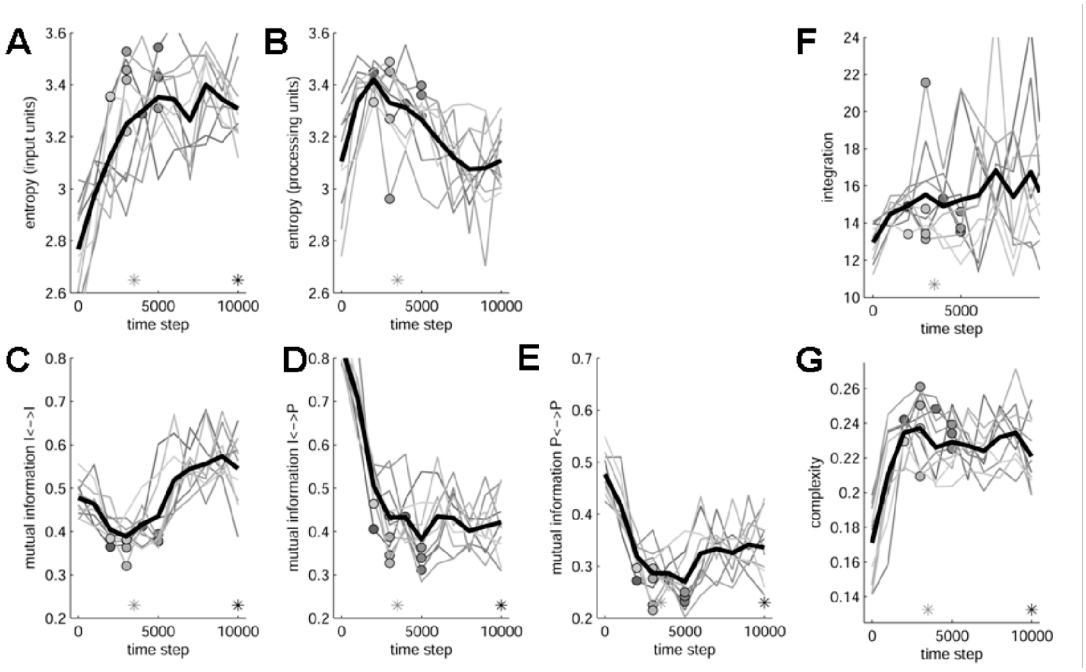

Figure 4: Information theoretic network metrics. (A,B) Entropy (C,D,E) Mutual information (F) Integration (G) Complexity.

Entropy. I neurons undergo a significant and persistent increase of entropy over evolution (Fig. 4A), possibly a result of increased population density in Polyworld. Entropy is elevated in P neurons at the point maximal population is reached, but decreases afterwards (Fig. 4B).

Mutual Information. For all binary relationships between I and P neurons, mutual information tends to decrease (Fig. 4C,D,E). This trend is especially pronounced for P neurons, indicative of increased functional specialization and de-correlation of their activity patterns.

Integration. Integration shows no clear trend across evolutionary stages. There is a slight, but possibly spurious, trend towards increased values, suggesting a weak tendency towards global integration (Fig. 4F). This trend appears to be largely offset by simultaneous decreases in mutual information indicating higher local specialization among I and P neurons.

Complexity. Both of these tendencies—increased functional specialization as expressed in decreased mutual information and a possible modest increase in global integration—are reflected in a statistically significant increase of complexity (Fig. 4G). The observed increase is modest, possibly reflecting the relative simplicity of the ecological environment and the fitness function.

Discussion

Within the confines of this artificial life system, with simple behavioral requirements for agents to thrive—survive, forage, and reproduce—we have observed and quantified consistent positive trends in structural elaboration and learning, that produce a statistically significant increase in neural complexity as a result of evolution.

What causes complexity to increase rather than decrease over time? Complexity, as measured in the present set of experiments, reflects the balance and co-existence of specialization and integration within a given neural architecture. Previous studies have indicated that complexity grows as neural systems become more effective in generating and integrating information (Sporns et al. 2000) and more interactive with their physical environments (Lungarella et al. 2005). We suggest that similar factors underlie the observed growth of neural complexity in Polyworld. Driven by a broadly defined fitness function that promotes a variety of behavioral traits and increased population density, neural systems in Polyworld are exposed to progressively higher rates of sensory input and must process this information to generate coherent behaviors. These challenges are met by the emergence of more structurally elaborate and more plastic networks, whose activity exhibits more differentiated as well as integrated dynamics as measured by complexity. We believe the implications of these trends may be profound.

First, the establishment of a general upward trend in complexity as a result of the action of evolution on biologically modeled nervous systems and behaviors suggests that such a trend may be present in natural biological systems. This particular simulation environment can probably be thought of as a single ecological niche, in which we have now observed a growth in complexity up until such time as the niche is fully exploited—until the population has reached a maximum and the individuals' expressed behaviors fully satisfy the only demands placed on them. We speculate that a well-formed measure of complexity applied to a biological species first occupying a new niche might exhibit similar growth. Then, given that all niches are not created equal, and that in more complex environments agent behaviors may result in additional niche creation, it is not difficult to imagine the observed growth in complexity extending to multiple niches and ecologies as a whole.

Second, the existence of trends in structural elaboration and learning capable of supporting a positive trend in complexity suggests that exploratory artificial life experiments may be able to first measure their progress and then continue to expand their horizons along this complexity scale. The earliest stated goals for Polyworld were to evolve its way up an intelligence spectrum, from the simplest levels of behavior to the most complex, adopting a "computational *Aplysia*" as a laudable first goal. We find the existence of a quantifiable, positive trend in evolved complexity in Polyworld encouraging.

Third, the existence of a much sought after "ruler", for quantitatively assessing complexity in artificial nervous systems, puts us in a much better position to expand our computational ecologies. There are an almost unlimited number of possible extensions to the existing world and agent models, and it appears we may finally have a quantitative scale with which to assess the benefits accrued from these extensions.

Also, the steady increase in Hebbian learning rate suggests that evolution strongly favors learning over innate behaviors in this system. Since we expect learning to play a major role in any fundamentally intelligent behavior, selection for learning might be seen as a required precursor to the evolution of intelligence, which requirement seems to have been met.

Future Directions

Having verified that evolution in Polyworld produces a positive trend in our information-theoretic measure of complexity, we are now in a position to apply this "ruler" to assess the benefits of changes to the system. We predict that increases in the complexity (now in the common sense of richness, variability, and useful predictability) of the environment will directly increase evolved agents' neural (information-theoretic) complexity.

In addition, both authors independently came to the realization, some years ago, that with a suitable measure of complexity in hand, complexity itself might be used

directly as a most interesting and effective fitness function in evolutionary simulations—a hypothesis for which there is now some evidence (Sporns and Lungarella 2006). To this end we intend to build the information-theoretic complexity calculations, currently performed externally, directly into Polyworld, and study the course of evolution in a computational ecology specifically designed to optimize for neural complexity.

Acknowledgements

We thank Nicole Di Camillo for useful discussions and comments, and Justin Post and Teresa Pegors for their early efforts at deploying these data analysis techniques on Polyworld data.

References

Adami, C., Ofria, C., and Collier, T. 2000. Evolution of biological complexity. PNAS 97(9):4463-4468.

Adami, C. 2002. What is complexity? BioEssays 24:1085-1094.

Bedau, M.A., Snyder, E., Brown, C.T., and Packard, N.H. 1997. A Comparison of Evolutionary Activity in Artificial Evolving Systems and in the Biosphere, in Proceedings of the Fourth European Conference on Artificial Life, 125-135. Cambridge, MA, MIT Press.

Bonner, J.T. 1988. The Evolution of Complexity by Means of Natural Selection. Princeton, NJ, Princeton Univ. Press.

Carroll, S.B. 2001. Chance and necessity: the evolution of morphological complexity and diversity. Nature 409:1102-1109.

Cover, T.M. and Thomas, J.A. 1991. Elements of Information Theory. New York, Wiley & Sons.

Crutchfield, J.P. and Young, K. 1990. Computation at the onset of chaos, in Complexity, Entropy, and Physics of Information, ed. Zurek, W. Reading, MA, Addison-Wesley.

Darwin, C. 1871. The Descent of Man, and Selection in Relation to Sex. Princeton, NJ, Princeton Univ. Press (1981).

Gould S.J. 1994. The Evolution of Life on Earth. Scientific American 271(4): 62-69.

Gray, L. 2003. A Mathematician Looks at Wolfram's New Kind of Science, Notices of the American Mathematical Society, 50(2):200-211.

Griffith, V. and Yaeger, L.S. 2006. Ideal Free Distribution in Agents with Evolved Neural Architectures, in Rocha, L. et al. eds. Artificial Life X. Cambridge, MA: MIT Press (this volume).

Heylighen, F. 1999. The Growth of Structural and Functional Complexity during Evolution, in The Evolution of Complexity: The Violet Book of `Einstein Meets Magritte'. eds. Heylighen, F.; Bollen, J.; and Riegler, A. Springer.

Langton, C.G. 1990. Computation at the Edge of Chaos: Phase Transitions and Emergent Computation, in Emergent Computation, 12-37. ed. Stephanie Forrest, North Holland.

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.

McShea, D.W. 1996. Metazoan complexity and evolution: is there a trend? Evolution 50:477–492.

Mitchell, M., Hraber1, P.T., and Crutchfield, J.P. 1993. Revisiting the Edge of Chaos: Evolving Cellular Automata to Perform Computations, Complex Systems 7:89-130.

Packard, N.H. 1988. Adaptation toward the edge of chaos. In Kelso, J.A.S., Mandell, A.J., and Shlesinger, M.F., eds. Dynamic Patterns in Complex Systems, 293–301, Singapore, World Scientific.

Papoulis, A. 1990. Probability and Statistics. Upper Saddle River, NJ, Prentice-Hall International Editions.

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, 109-118. Heidelberg, Springer-Verlag.

Shalizi, C.R. 2003. Optimal Nonlinear Prediction of Random Fields on Networks, in Discrete Mathematics and Theoretical Computer Science, AB(DMCS) 11-30 and http://arxiv.org/abs/math.PR/0305160.

Sporns, O., Tononi, G. and Edelman, G.M. 2000. Theoretical neuroanatomy: Relating anatomical and functional connectivity in graphs and cortical connection matrices. Cerebral Cortex, 10:127-141.

Sporns, O. and Lungarella, M. 2006. Evolving Coordinated Behavior by Maximizing Information Structure. In Rocha, L. et al. eds. Artificial Life X. Cambridge, MA: MIT Press (this volume).

Tononi, G., Sporns, O. and Edelman, G.M. (1994). A measure for brain complexity: Relating functional segregation and integration in the nervous system. Proc. Nat. Acad. Sci. USA, 91:5033-5037.

Wolfram, S. 1984. Universality and complexity in cellular automata, Physica D 10:1-35.

Yaeger, L.S. 1994. Computational Genetics, Physiology, Metabolism, Neural Systems, Learning, Vision, and Behavior, in Proceedings of the Artificial Life III Conference, 263-298. ed. Langton, C. G. Addison-Wesley.