Functional and Structural Topologies in Evolved Neural Networks

Joseph T. Lizier[1,2,⋆], Mahendra Piraveenan[1,2], Dany Pradhana[1],
Mikhail Prokopenko[1], and Larry S. Yaeger[3,⋆⋆]

[1] CSIRO Information and Communications Technology Centre, Locked Bag 17,
North Ryde, NSW 1670, Australia
[2] School of Information Technologies, The University of Sydney, NSW 2006, Australia
[3] School of Informatics, Indiana University, 919 E. 10th St., Bloomington, IN 47408,
USA
`mahendra.piraveenan@csiro.au`

Abstract. The topic of evolutionary trends in complexity has drawn much controversy in the artificial life community. Rather than investigate the evolution of overall complexity, here we investigate the evolution of topology of networks in the Polyworld artificial life system. Our investigation encompasses both the actual structure of neural networks of agents in this system, and logical or functional networks inferred from statistical dependencies between nodes in the networks. We find interesting trends across several topological measures, which together imply a trend of more integrated activity across the networks (with the networks taking on a more "small-world" character) with evolutionary time.

1 Introduction

The nature of evolutionary trends in complexity has been subject to much debate [1], with interest surrounding whether the evolutionary growth in complexity of organisms in the natural world is the outcome of natural selection or some sort of random walk [2, 3]. Indeed, this question has been explored in artificial life systems: e.g. previous work with Polyworld has demonstrated that evolution can and does select for increased complexity in a driven fashion in *some* circumstances, but also selects for complexity stability under other conditions [4, 5].

Here, our interest lies not so much in the evolution of (any particular measure of overall) complexity, but rather the manner in which the *topology* of neural networks adapt under evolutionary pressure. Specifically, we investigate the

---

[⋆] Authors are in alphabetical order.

[⋆⋆] The authors thank the sponsors of the Guided Self-Organisation Workshop 2008 (GSO-2008), who partially supported this work: the Australian Research Council's Complex Open Systems Research Network (COSNet) and Research Network in Enterprise Information Infrastructure (EII), The University of Sydney, and CSIRO Complex Systems Science and ICT Centre.

topology of neural networks of agents in the Polyworld artificial life system. We examine both the actual structure of these networks, and their logical structure.

The logical structure of the neural networks is explored by inferring functional networks [6, 7] from statistical dependencies between the time series of each node in the underlying structural network. Here, we use mutual information [8] and transfer entropy [9] to measure the statistical dependencies between the neurons. We then examine the trends in several measures of the topology of the structural and functional networks with respect to evolutionary time: in particular, we measure the assortativity, modularity, clustering coefficient and closeness of the networks. We find several interesting trends in the topologies, with the trends in the structural and transfer entropy-based functional networks being most similar. These networks become more non-assortative, less modular but more clustered, and adopt shorter average path lengths with evolutionary time. These trends are significant in that they imply the networks are taking on a more "small-world" [10] character over evolutionary time.

2 Polyworld

Polyworld [11] is a computational ecology evolving populations of haploid agents, each using a suite of primitive behaviors (move, turn, eat, mate, attack, light, focus) under continuous control of an Artificial Neural Network (ANN) employing summing and squashing neurons with synapses that adapt via Hebbian learning. The wiring diagram of the ANN is encoded in the organism's genome, via a statistical description of the number of neural groups of excitatory and inhibitory neurons, synaptic connection densities, ordered-ness of connections, 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. The agent morphologies are simple and fixed, but agents' interactions with the world and each other are fairly complex, as they replenish energy by seeking out and consuming food or by killing and eating other agents. They reproduce when two collocated agents simultaneously express their mating behaviors, using a number of crossover points and a mutation rate that are also contained in the parental genomes [11].

Bounds on the agent population, both high and low, are maintained by altering the energy consumption of the agents (as in [5]). As the population approaches the upper bound, the amount of energy depleted by all agent behaviors, including neural activity, is increased in a continuous fashion. Reciprocally, as the agent population approaches the minimum, energy depletion is decreased, and agent lifespans may be artificially extended.

The simulation is initially seeded with a uniform population of agents that have the minimum number of neural groups and a nearly minimal number of neurons and synapses. While predisposed to some potentially beneficial behaviors, such as running towards food (green) and away from aggression (red; see [11] for details on color use in Polyworld), these seed organisms are not a viable species. I.e., without evolution they cannot sustain their numbers through their reproductive behaviors and will inevitably die out.

As simulations progress both the structural architecture of the ANNs and the activation of every neuron at every time step are recorded for every agent. Here we use these neural activation recordings to determine functional networks for each agent and compare functional network characteristics to the underlying structural network characteristics.

3 Inferring Functional Networks

Two remote neural nodes are defined to be functionally connected where they exhibit statistical dependence in time [6, 7]. The *nodes* considered could be voxels in BOLD recordings (e.g. [7]), or neurons in an artificial neural network (as are used here). A functional network is then formed from a set of functional connections. Inferring functional networks from time-series of node states therefore involves two distinct steps: i. making some measure of the statistical dependence or closeness between each node pair, then ii. deciding whether each closeness value should constitute a link between the node pair. The closeness measure and the inferred links can be either directional or undirectional.

Functional networks may be used to infer the underlying structural network where this is unknown. More importantly, functional networks provide insight into the *logical* structure of the network and how this changes as a function of network activity (regardless of whether the underlying structure is known).

In this work, we use information-theoretical measures [8] for the closeness of each pair X and Y. The **mutual information** between X and Y measures the average reduction in uncertainty about x (or entropy H of x) that results from learning the value of y, or vice versa:

$$I(X;Y) = \sum_{x,y} p(x,y) \log_2 \frac{p(x,y)}{p(x)p(y)}. \tag{1}$$

In this way, $I(X;Y)$ is a symmetric measure of the common information between X and Y. Though it has been previously used to measure directed information transfer from one variable to another, this is not valid: it is a symmetric measure of *statically shared* information (which is useful in its own right).

Alternatively, the **transfer entropy** [9] was introduced as a *directed* measure of *dynamic* information transfer from one variable to another. It quantifies the information provided by a source node about a destination's next state that was not contained in the past of the destination. Specifically, the transfer entropy from a source node Y to a destination X is the mutual information between the previous state of the source y_n and the next state of the destination x_{n+1}, *conditioned* on the past k states of the destination $x_n^{(k)}$:

$$T_{Y \to X}(k) = \sum_{x_{n+1}, x_n^{(k)}, y_n} p(x_{n+1}, x_n^{(k)}, y_n) \log_2 \frac{p(x_{n+1} | x_n^{(k)}, y_n)}{p(x_{n+1} | x_n^{(k)})}. \tag{2}$$

The transfer entropy may be measured for any two time series X and Y and is always a valid measure of the predictive gain from the source, but only represents physical information transfer when measured on a causal link [12].

Here, we compute functional networks for each agent from the Polyworld simulation using both mutual information and transfer entropy as separate measures of closeness. The continuous activation levels are first discretised in four levels, and a history length $k = 1$ is used for the transfer entropy (this renders it more towards an inference of causal effect than information transfer [13, 12]).

Several options are then available for deciding whether each pair of areas should be considered functionally connected based on their closeness. One could assign links to a given number or percentage of pairs based on the largest closeness values, or could use an approach based on the statistical significance of the closeness measure, e.g. [14]. Here, the number of functional links was designed to match the proportion of links in the underlying structural network, and the largest such closeness values were assigned links. A (directed) link exists in the structural network between two neurons where the source neuron is an input to the target neuron. We consider both processing and input neurons in the functional network.

4 Network Topological Measures

Analysis of the *topology* of functional networks provides useful information about the dynamic behaviour of the network [7, 14]. In this section, we introduce the measures of topology used to analyse the functional networks here. All were calculated using [15].

Assortativity is the tendency observed in networks where nodes mostly connect with similar nodes. Typically, this similarity is interpreted in terms of degrees of nodes. Assortativity has been formally defined as a correlation function of excess degree distributions and link distribution of a network [16, 17]. The concepts of degree distribution $p(k)$ and excess degree distribution $q(k)$ for undirected networks are well known [17]. Given $q(k)$, one can introduce the quantity $e_{j,k}$ as the joint probability distribution of the remaining degrees of the two nodes at either end of a randomly chosen link. Given these distributions, the assortativity of an undirected network is defined as:

$$r = \frac{1}{\sigma_q^2} \left[\sum_{jk} jk \left(e_{j,k} - q\left(j\right) q\left(k\right) \right) \right], \tag{3}$$

where σ_q is the standard deviation of $q(k)$. Assortativity distributions can be constructed by considering the *local assortativity* of all nodes in a network [18].

Closeness centrality of a node v is defined as the mean geodesic distance (shortest path length) between the node and all other nodes in the network [19]. (Sometimes the quantity is inverted so that the nodes which are 'most central' to the network G would get higher values). Closeness centrality is formally defined as $C^C(v) = \sum d_G(v,t)$ where $v \neq t$ and $d_G(v,t)$ is the shortest path distance between nodes v and t.

Network **modularity** is the extent to which a network can be separated into independent sub-networks. Formally [20], modularity quantifies the fraction of links that are within the respective modules compared to all links in a network. [20] introduces an algorithm which can partition a network into k modules and measure the partition's modularity Q. The measure uses the concept that a good partition of a network should have a lot of within-module links and a very small number of between-module links. The modularity can be written as:

$$Q = \sum\nolimits_{s=1}^{k} \left[\frac{l_s}{L} - \left(\frac{d_s}{2L} \right)^2 \right], \tag{4}$$

where k is the number of modules, L is the number of links in the network, l_s is the number of links between nodes in module s, and d_s is the sum of degrees of nodes in module s. To avoid getting a single module in all cases, this measure imposes $Q = 0$ if all nodes are in the same module or nodes are placed randomly into modules.

The **clustering coefficient** of a node characterizes the density of links in the environment closest to a vertex. Formally, the clustering coefficient C of a node is the ratio between the total number y of links connecting its neighbours and the total number of all possible links between all these z nearest neighbours [21]: $C = 2y/\left(z\left(z-1\right)\right)$. The clustering coefficient $\overline{C}$ for a network is the average C over all nodes.

5 Results and Discussion

We constructed the functional networks for each agent, and evaluated each measure of network topology on these and the underlying structural networks (which had between 13 and 159 neurons, and 52 on average). We then averaged each measure over sets of 100 sequential agents ordered by birth. The results are plotted with respect to evolutionary time in Fig. 1. Clearly, all measures reach a relatively steady state within 5000 – 12000 steps in evolutionary time. This aligns with previous studies of trends in the complexity of the neural networks in Polyworld [5] where the complexity is driven upwards over the initial 5000 or so steps of evolution before the agents find a "good enough" solution. At this point the drive for evolutionary change somewhat stagnates, as is reflected in the steady state of the measures here.

In general, the transfer entropy-inferred functional networks show similar trends to the structural networks across all measures. Interestingly, the transfer entropy-inferred functional networks had a slightly smaller overlap (mean $17.6 \pm 0.1\%$) with the underlying structural networks than the mutual information-inferred functional networks (mean $19.1 \pm 0.1\%$). It is possible that the transfer entropy performs better at inferring the general interaction structure between modules or regions in the structural network (thereby capturing the general topological trends) without necessarily inferring the precise links any better.

As shown in Fig. 1(a), the structural networks tend to exhibit a negative assortativity: this is not surprising as it is a known general characteristic of biological networks evolved under external pressure [22]. This is because negative

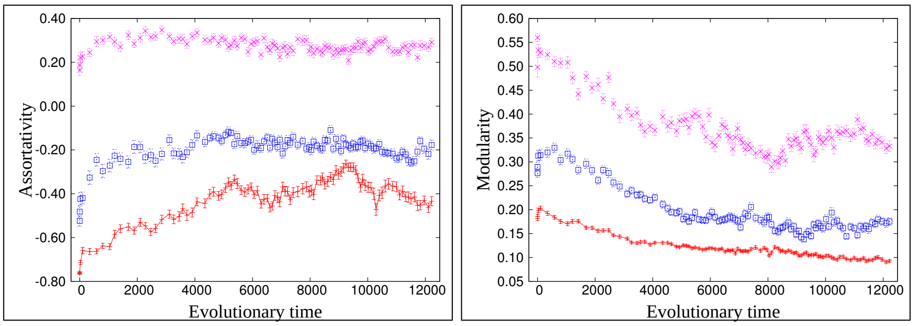

(a) Assortativity (b) Modularity

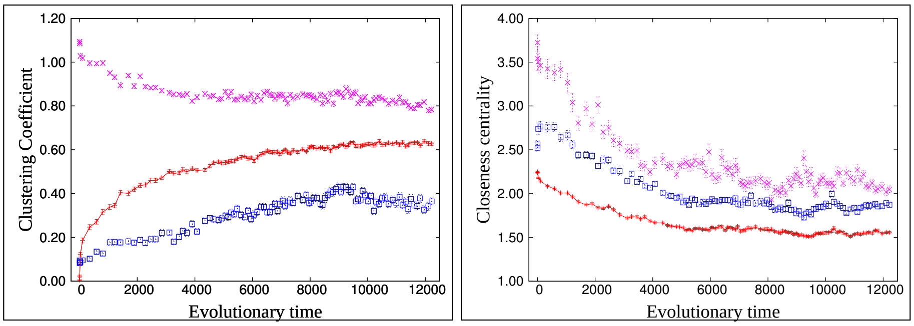

(c) Clustering coefficient (d) Closeness Centrality

Fig. 1. Trends in structural and functional networks versus evolutionary time. Measures are plotted for structural networks (red line), mutual information-inferred functional networks (violet ×), and transfer entropy-inferred functional networks (blue ☐). Error bars indicate the standard error of the mean.

assortativity supports connectivity between diverse elements in the network, an important feature for producing complex behaviour. Unsurprisingly also, the mutual information-inferred networks exhibit positive assortativity (since mutual information is maximised for similar elements), while the transfer entropy-inferred networks exhibit negative assortativity (since transfer entropy is minimised for similar elements). More interestingly, the structural and transfer entropy-inferred networks become more neutrally assortative over time (i.e. less negatively assortative). While this may seem surprising, it is possibly an artifact of the elements in the network becoming more closely coupled as they evolve and therefore become more similar, or perhaps reflects the increased clustering occurring over evolutionary time.

Fig. 1(b) and Fig. 1(c) show that the structural and transfer entropy-inferred networks become less modular but more clustered as they evolve. This is not a contradiction: it indicates that the boundaries between modules are becoming blurred with previously separated modules becoming more strongly clustered both *within* themselves and *across* each other (i.e. finding the right balance between functional integration and segregation to give rise to complex behaviour).

The mutual information-inferred networks however exhibit a decrease in clustering coefficient. Again, this seems to be a relic of the mutual information measure being maximised for similar elements: stronger coupling across clusters in the underlying network is likely to diversify the activity of previously similar nodes, thereby reducing clustering in this functional network.

Finally, Fig. 1(d) shows that the closeness centrality is reduced with evolutionary time for all networks. Given the previous results, this is unsurprising as all imply diversification of connectivity across the network with evolutionary time. In fact, taken together these results (in particular the higher clustering and lower shortest path lengths) suggest that the networks are becoming more small-world [10] with evolutionary time. Again this is unsurprising but significant, since the same effect is observed in many natural systems (including biological cortical networks and networks optimised for complexity [23], as well as functional networks inferred from neural networks in [14]) due to the advantages bestowed by this property. Importantly though, recall that all measures reach a steady state here: the neural networks do not continually improve on these desirable features, but stop developing once a good enough solution is found.

6 Conclusion

We have measured functional networks to represent the logical activity of neural networks of agents in the Polyworld artificial life system. Topological analysis of these functional networks, and the underlying structural networks, revealed clear trends with evolutionary time. The structure and activity in the networks becomes more integrated over time, as may be expected in the evolution of complex distributed processes. In particular, both the structural and functional networks take on more of a small-world character as the evolution progresses.

Our results also showed interesting differences between the use of mutual information and transfer entropy in inferring functional networks. The transfer entropy-inferred functional networks have topological trends more similar to those of the underlying structural networks, and also provided more intuitive insights into network activity.

In extending this work, it would be desirable to evaluate the statistical significance of the trends observed here. One method for doing this would be to contrast the results here (where evolution is driven by genetic mixing) with those produced by passive genetic drift (along the same lines as the comparison of trends in complexity in [5]).

References

1. Bedau, M.A.: The evolution of complexity. In Barberousse, A., Morange, M., Pradeu, T., eds.: Mapping the Future of Biology. Volume 266 of Boston Studies In The Philosophy Of Science. Springer, Netherlands (2009) 111–130
2. Gould, S.J.: The evolution of life on earth. Scientific American **271**(4) (1994) 62–69

3. Maynard Smith, J.: Time in the evolutionary process. Studium Generale **23** (1970) 266–272
4. Yaeger, L., Sporns, O.: Evolution of neural structure and complexity in a computational ecology. In Rocha, L.M., Yaeger, L.S., Bedau, M.A., Floeano, D., Goldstone, R.L., Vespignani, A., eds.: Proceedings of the Tenth International Conference on Simulation and Synthesis of Living Systems (ALifeX), Bloomington, Indiana, USA, MIT Press (2006) 330–336
5. Yaeger, L., Griffith, V., Sporns, O.: Passive and driven trends in the evolution of complexity. In Bullock, S., Noble, J., Watson, R., Bedau, M.A., eds.: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems (ALifeXI), Winchester, UK, Cambridge, MA, USA, MIT Press (2008) 725–732
6. Friston, K.J.: Functional and effective connectivity in neuroimaging: A synthesis. Human Brain Mapping **2** (1994) 56–78
7. Honey, C.J., Kotter, R., Breakspear, M., Sporns, O.: Network structure of cerebral cortex shapes functional connectivity on multiple time scales. Proceedings of the National Academy of Sciences **104**(24) (2007) 10240–10245
8. MacKay, D.J.: Information Theory, Inference, and Learning Algorithms. Cambridge University Press, Cambridge (2003)
9. Schreiber, T.: Measuring information transfer. Phys. Rev. Lett. **85**(2) (2000) 461–464
10. Watts, D.J., Strogatz, S.: Collective dynamics of small-world networks. Nature **393** (1998) 440–442
11. Yaeger, L.S.: Computational genetics, physiology, metabolism, neural systems, learning, vision and behaviour or polyworld: Life in a new context. In Langton, C.G., ed.: Proceedings of the Artificial Life III Conference, Santa Fe, NM, USA, Addison-Wesley (1994) 263–298
12. Lizier, J.T., Prokopenko, M.: Differentiating information transfer and causal effect (2009) arXiv:0812.4373.
13. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: Local information transfer as a spatiotemporal filter for complex systems. Phys. Rev. E **77**(2) (2008) 026110
14. Bettencourt, L.M.A., Stephens, G.J., Ham, M.I., Gross, G.W.: Functional structure of cortical neuronal networks grown in vitro. Phys. Rev. E **75**(2) (2007) 021915
15. Sporns, O., Rubinov, M., Kötter, R.: Brain connectivity toolbox (2009) http://www.brain-connectivity-toolbox.net/.
16. Newman, M.: Assortative mixing in networks. Phy. Rev. Lett **89**(20) (2002) 208701
17. Newman, M.: Mixing patterns in networks. Phy. Rev. E **67**(2) (2003) 026126
18. Piraveenan, M., Prokopenko, M., Zomaya, A.Y.: Local assortativeness in scale free networks. Euro. Phys. Lett **89**(2) (2008) 28002
19. Shimbel, A.: Structural parameters of communication networks. Bulletin of Mathematical Biology **15**(4) (1953) 501–507
20. Alon, U.: An Introduction to Systems Biology: Design Principles of Biological Circuits (Chapman & Hall/Crc Mathematical and Computational Biology Series). Chapman & Hall/CRC (July 2006)
21. Dorogovtsev, S., Mendes, J.: Evolution of Networks: From Biological Nets to the Internet and WWW. Oxford University Press, Oxford (January 2003)
22. Solé, R.V., Valverde, S.: Information theory of complex networks: Onevolution and architectural constraints. In Ben-Naim, E., Frauenfelder, H., Toroczkai, Z., eds.: Complex Networks. Volume 650 of Lecture Notes in Physics. Springer, Berlin / Heidelberg (2004) 189–207
23. Sporns, O., Kötter, R.: Motifs in brain networks. PLoS Biology **2**(11) (2004) e369