Individual and Global Adaptation in Networks

Richard A. Watson, Rob M. Mills, Simon T. Powers,
Adam Davies, Chris Buckley[1], Chris Cox, Adam Jackson

Natural Systems Group, University of Southampton, U.K.
[1]Informatics, Sussex University, UK.
raw@ecs.soton.ac.uk

Abstract

The structure of complex biological and socio-economic networks affects the selective pressures or behavioural incentives of components in that network, and reflexively, the evolution/behaviour of individuals in those networks changes the structure of such networks over time. Such 'adaptive networks' underlie how gene-regulation networks evolve, how ecological networks self-organise, and how networks of strategic agents co-create social organisations. Although such domains are different in the details, they can each be characterised as networks of self-interested agents where agents alter network connections in the direction that increases their individual utility. Recent work shows that such dynamics are equivalent to associative learning, well-understood in the context of neural networks. Associative learning in neural substrates is the result of mandated learning rules (e.g. Hebbian learning), but in networks of autonomous agents 'associative induction' occurs as a result of local individual incentives to alter connections. Using results from a number of recent studies, here we review the theoretical principles that can be transferred between disciplines as a result of this isomorphism, and the implications for the organisation of genetic, social and ecological networks.

Neural (Adaptive) Networks

The Hopfield network [6] is a well-understood example of a dynamical system based on pairwise interactions (Eq. 1). It has provided a vehicle for studying dynamical behaviour across many disciplines from neural networks, to spin-glass models, to ecosystems. For example, it can be relabelled to represent strategic choices of agents rather than activation responses of neurons. Specifically, each agent in a network, repeatedly, one at a time in random order, adopts one of two discrete behaviours or states, $s_i = \pm 1$ (e.g. trait-a/trait-A, trade-x/trade-y, vote-A/vote-B), so as to maximise an individual pay-off, fitness or utility, u_i, which is a weighted sum of interactions between the state it adopts and the states of other agents in the network,

$$u_i = \sum_{j}^{N} \omega_{ij} s_i s_j \qquad (1)$$

where $\omega_{ij} = (-1,1)$ is the current interaction coefficient or connection between agent i and agent j. In principle, an agent may have interactions with all other agents, but non-connected agents may be represented by $\omega_{ij} = 0$.

The most well-known capability of the Hopfield network is its use as a content addressable memory [6]. If the weights are appropriately configured, the attractors of the system can be used to store patterns of activation (state configurations) that can be recalled from partial or corrupted stimuli simply by running the network from the stimulus pattern to a local attractor. This exhibits the property that the network will 'fill-in' or repair missing or corrupted data in the input pattern. A second well-known use of the Hopfield network is as an optimisation method; the weights of the network are used to represent the constraints of an optimisation problem and running the network from a random initial condition finds a locally optimal solution to those constraints [7]. In agent terms, the latter behaviour corresponds to finding a local Nash equilibrium in the network of mutually constrained coordination ($\omega_{ij}>0$)/anti-coordination ($\omega_{ij}<0$) games.

Associative learning

For the memory capability, local learning rules can be used to train a network to store particular patterns without hand-designing the weights. Hebbian learning [5] is simply the idea that learning occurs by changes to the connections between neurons, and Hebb's rule (or the delta rule) states that the change in a weight is proportional to the product of the states it connects:

$$\Delta\omega_{ij} = \delta\, s_i s_j . \qquad (2).$$

where $\delta > 0$ is a learning rate. To store a particular pattern, the states are set to the pattern and the weights are changed using this rule. Likewise, to store a set of patterns, the weights are adjusted incrementally for each pattern in turn repeatedly. In this case, when a subset of patterns in the training set are very similar (not orthogonal), only one memory (attractor state) may be learned for them. In storage or memory terms where the intention is to recall all patterns faithfully, this is undesirable. But, in other cases this same characteristic is a desirable form of simple generalization – enabling the network to 'recognise' the class of pattern and return the exemplar of the class, rather than merely returning the most-similar training pattern. It is important to note that Hebbian learning is an unsupervised learning method (weight changes are not assessed with respect to an external performance metric), and use only local information, i.e. the change in a weight is a function only of the two states that it connects.

Thus the Hopfield network has non-trivial (but well-understood) state dynamics with interesting properties, and likewise, well-understood topological dynamics. But, in these applications, these dynamics are separate – i.e., the learning phase precedes a distinct recall phase. Some work investigates coupled state-topology coevolution [4] in the context of such networks [17] but for different purposes to those studied here.

‘Self-modelling’ Systems

What happens when a neural network ‘learns a model of its own behaviour’? Our recent work shows that, under certain circumstances, when state and topology processes occur simultaneously in the same network (but on different timescales) this can have interesting consequences for the optimisation capabilities of the Hopfield network. Specifically, unsupervised distributed learning improves the ability of the network to solve constraint optimisation problems [25].

To understand this behaviour, first note that the effect of Hebbian learning is to enlarge the basin of attraction (or create one de novo) for the current state configuration – hence, it increases the number of initial conditions that will reach this configuration – which is what is meant by a memory. Second, in systems built out of many low-order interactions (e.g., pairwise interactions or two-player games) there are limits on the steepness of slopes in the energy landscape (a single state change affects at most N out of N^2 constraints) and this ensures, on average, positive correlation between the depth of attractors and the width of their basins [25]. This implies that the globally optimal solution to a constraint problem, encoded in the weights of the network, will probably have the largest basin of attraction for the network dynamics – even though its basin might be only a tiny fraction of the entire configuration space (i.e., it is a difficult optimisation problem for local search). There also necessarily exists significant commonality between the combinations of states that satisfy (some) constraints at local optima, and the combinations of states that satisfy (more) constraints at global optima. These observations mean that learning tends to increase the reliability with which the system is attracted to good quality solutions, and together with the generalisation ability of associative learning, a network can generalise over a set of local optima to enlarge the basin of the global optimum even before it is sampled, decreasing the time to first hit of the global optimum.

Given that a problem can be encoded in the weights of the network (as a MAXSAT problem [25]), the conditions necessary for this result are quite simple. Specifically, if the state dynamics are repeatedly perturbed (re-initialised at random) and allowed to settle to a locally optimal configuration, such that the weights change slowly whilst the state dynamics visit a distribution of locally optimal state configurations, then these locally optimal solutions provide the training samples necessary for associative generalisation to learn the location of the global optimum [25].

In energy minimisation terms, rather than optimisation terms, the system finds a topology that enables it to minimise the energy of the state variables more reliably and more completely. In other words, the network represents a dynamical system with a sort of second-order energy minimisation ability.

Equivalence of Associative Learning and Self-Interested Changes to Connections

Above we described the effect of state-topology coevolution in the Hopfield network on the assumption that weight changes follow a mandated Hebbian principle. But what happens if we cannot mandate such a learning rule – what happens if autonomous agents in an adaptive network change the connections of a network to suit their own interests? Our work shows that, in fact, the result is exactly the same. That is, if agents in a network alter the strength of connections in the direction that increases their utility (given the state configuration they currently experience) they will necessarily alter that connection in the Hebbian direction [28]. Actually, the mathematics of this is trivial (3) –

$$u_i' > u_i \Rightarrow \sum_{j}^{N} \omega_{ij}' s_i s_j > \sum_{j}^{N} \omega_{ij} s_i s_j$$

$$\Rightarrow \begin{cases} \omega_{ij}' > \omega_{ij} \quad if \quad s_i s_j > 0, & or \\ \omega_{ij}' < \omega_{ij} \quad if \quad s_i s_j < 0. & \end{cases} \qquad (3)$$

And with some reflection, the reason becomes clear. When an agent changes connections selfishly they must change the connection to make themselves better off given the current state configuration. This directly makes the current state configuration more desirable to the agent, and in some cases, preferred over other state configurations where previously it was not. Thus, when an agent increases the desirability of a state configuration by changing connection strengths, it necessarily has the side-effect of increasing the basin of attraction for that state configuration. Increasing the utility, or decreasing the energy, of a state configuration, necessarily widens its basin of attraction, which is why Hebb’s rule has to take the form it does, and why selfish changes to connections must have the side-effect of creating an associative memory.

In fact, it is not even necessary that the causal process that changes the connections is intelligent or the result of a facultative agent ‘intending’ to increase utility; the same effect is observed in a purely mechanical system of particles joined by springs. If the springs are a little plastic, if their natural length deforms slowly under stress (i.e., they are not perfect springs), the result is a mechanical system that forms an associative memory of the patterns of stress that it has been exposed to in the past [30]. To acknowledge the equivalence with associative learning, but to emphasise that the mechanism is not centrally mandated, we call this process *associative induction*.

The universality of this mechanism is indicated in the observation that Hebb’s rule is simply the equation of motion for ω given that its behaviour is governed by the *same* energy function (or utility function) as the states.

This equivalence with Hebbian learning and the simplicity of the mechanism (being distributed and unsupervised, without global feedback on total utility, and based on purely self-interested incentives) implies that it may be relevant to a broad class of adaptive networks. This suggests that many different types of adaptive networks in diverse domains may in principle be capable of the same learning, memory, and generalisation abilities well-understood in the context of neural networks – including the improvement in optimisation ability observed in self-modelling systems.

Thus when any adaptive network of this form is conditioned with the perturbation protocol mentioned above, i.e. repeatedly reset and allowed to relax to different local optima, it will – exactly as in the case of the self-modelling Hopfield network – learn to optimise better. From the point of view of local incentives versus global optimisation, this is rather interesting. It means that selfish agents, acting purely to optimise their individual utility, under these simple conditions, nonetheless restructure their interactions so that the global utility – sum of individual utilities – is optimised, or more generally, that the system finds more optimal configurations more reliably [28,25]. The equivalence also helps us understand the limitations of the processes, e.g., just as a high learning rate would cause a Hopfield network to overfit to training data, when the rate of change of connections is too high, a network will learn a model of the particular attractors that it happens to visit early in the process and fail to build a generalised model or find good solutions.

In dynamical system terms, we may imagine the original behaviour of the network simply as a ball rolling down hill – albeit a multi-dimensional ball on a multi-dimensional and rugged hill. This simply finds a local minimum. But as the system visits a distribution of local optima, and the connections are slowly deformed by the states to incrementally relieve the stresses produced by the state configurations at each local optimum visited, the ball 'gets better' at rolling downhill. That is, the learned weights sometimes create situations where the ball takes a trajectory that it couldn't previously have taken. From the ball's point of view, given the new weights, this is simply the obvious thing to do – the downhill direction. But with respect to the original energy function the new trajectory sometimes takes the ball uphill (with respect to the original energy function/objective function), but in a direction that subsequently enables it to find lower minima.

We have been investigating the implications of these findings in a number of different domains. In **gene regulation networks** [24] we find that selected changes to regulatory connections tend to canalise the locally optimal phenotype (forming a memory of it, so that it is produced more reliably) by increasing the co-regulation of traits that are selected together. And if this occurs over a distribution of locally optimal phenotypes (i.e. phenotypes resulting from different selective environments [16], or from different initial conditions in the same selective environment [24]) it enhances evolvability – enables evolution to find fitter phenotypes. This implies a formal resolution to the seemingly opposing notions of phenotypic robustness and evolvability [9]; and suggests formal principles for characterising evo-devo dynamics [20] and the relationship between proximal selective pressures and ultimate adaptive consequences.

In **social multi-agent systems**/games on networks [3], we find that selfish changes to connections habituate the system to the current Nash equilibrium by increasing the perceived utility of interactions that have been successful in the past. But if selfish agents modify connections over a distribution of local Nash (e.g., in a repeated social task), this effects an equilibrium selection process that causes the system to find (non-local) equilibria of higher total-utility. These works enable us to relate simple local strategies such as reinforcement learning, or habituation, with system-level maximisation of social welfare. Ongoing work in this area also conceives changes to weights as changes to a 'meta-game', i.e., where agents' strategies effectively, but subtly, change the game they are playing [8].

In **ecological models** [31] (non-trophic inter-specific fitness dependencies, such as from resource competition), we find that if species in an ecosystem can evolve ecological interactions or symbiotic relationships with other species then species that commonly co-occur evolve interactions that reinforce their co-occurrence in future. This has the effect of enlarging the attractor for the current ecological equilibrium, and if this occurs over a distribution of local ecological equilibria (e.g. under conditions of occasional ecological disturbances) it enlarges the basin for equilibria that are especially efficient in utilising available resources and thus attain higher biomass. Ongoing work in this area [26,27], together with the work on meta-games, is beginning to formally integrate with social evolution theory (kin selection and inclusive fitness [11]) via a process of 'social niche construction', i.e., selective pressures on individuals to change who they interact with and how much [18]. This has been investigated in simple structures (e.g. group size [19,21]) and here extends into network-based structures where agents change the specific membership of groups (not just the relatedness of group members).

Transformations in the Scale of Behaviour

The above investigations address the use of associative learning principles to understand how the structures of adaptive networks change under individual incentives, and how this changes alignment with total social welfare. These works involve relatively subtle processes of altering the system dynamics. However, we have also been investigating the use of more radical modes of altering dynamical processes based on associative learning processes that alter system structure. Specifically, we have developed optimisation algorithms in a class of algorithms we call 'multi-scale search' – search algorithms that use local search at one (primitive) scale to identify modules that can be used to rescale the variation operators (or 'coarse grain' the neighbourhood of the search space) and thereby reduce the effective search space, focussing intelligently on fit regions; this is applied recursively through successive scales of organisation [12] (also [14]).

This work, again, uses only distributed/local processes of association-building that are Hebbian. The result of these associations is that the system of variables effectively becomes a higher-level (lower-dimensional) model of its own behaviour, and again, associative generalisation means that this can find high-quality solutions faster. The fact that the association-building mechanism is Hebbian, and therefore in alignment with individual incentives, makes this optimisation method relevant to networks of self-interested agents. That is, the system is equivalent to a network of agents that form 'coalitions', enabling coordinated strategy choices and pooled utility [27,29]. These coalitions enable groups of agents to escape locally-optimal Nash equilibria, by changing many player-strategies in a coordinated fashion, and thereby escape to other Nash equilibria of higher total utility [29,31].

An individual-based simulation [31,27] shows that when individuals can evolve (generative [29]) associations with other species, they spontaneously implement the same algorithm. This, we suggest, provides a model of how the formation of new selective units in the major transitions in evolution [10,1,15] can be driven by bottom-up individual incentives, and yet alter evolutionary outcomes significantly and systematically in the direction of high total welfare.

This begins to suggest that there are additional algorithmic principles involved in natural evolution [23], and supports a view of evolution that improves its' own ability to create adaptive variants over time [9,16,22,2] – by creating new evolutionary units and scaling-up the adaptive process [10,1, 29,27,31,12]. This work also has potential to shed light on the self-organisation of social structure in socio-economic complex systems – perhaps even to provide a formal basis with which to understand the conditions where selfish agents co-create social contracts that maximise social welfare, and the conditions where they don't. Thus far, conceiving these varied systems as adaptive networks, and exchanging concepts with the field of connectionist learning where such dynamics are better understood has yielded significant insight about the relationship between individual and global adaptation in complex biological networks.

Acknowledgements Thanks to Paul Ryan and Jason Noble for conceptual discussion.

References

1. Buss, LW, (1987), *The Evolution of Individuality*, Princeton Press, New Jersey
2. Calcott, B. & Sterelny, K. (2011) "Introduction: A Dynamic View of Evolution" in *The Major Transitions in Evolution Revisited*, MIT Press.
3. Davies, A., Watson, R.A., Mills, R.M., Buckley, C.L., Noble, J. (2011) If You Can't Be With the One You Love, Love the One You're With: How Individual Habituation of Agent Interactions Improves Global Utility. *Artificial Life*. 17(3):167-81.
4. Gross, T. & Sayama, H. (2009) *Adaptive Networks. Theory, Models and Applications*. Springer-Verlag: Berlin.
5. Hebb, D.O. (1949). *The organization of behaviour*. New York: Wiley.
6. Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities, *PNAS USA*, 79 (8) 2554-2558.
7. Hopfield, J.J., & Tank, D.W. (1985). 'Neural' computation of decisions in optimization problems. *Biol Cybern* 52:141-152.
8. Jackson, A. (2011) "Social Dilemmas and Meta-Games", Summer (MSc) dissertation, Sept 2011, ECS, University of Southampton.
9. Kirchner, M. & Gerhart, J. (1998). Evolvability. *PNAS. USA*, 95:8420–8427
10. Maynard Smith, J. & Szathmary, E. (1995) *Major Transitions in Evolution*. W. H. Freeman.
11. Michod, R. E. & Hamilton, W. D. (1980) Coefficients of relatedness in sociobiology. *Nature* 288, 694 – 697.
12. Mills (2010) *How Micro-Evolution Can Guide Macro-Evolution*, PhD thesis, ECS, Southampton.
13. Mills, R., Watson, R. A. and Buckley, C. L. (2011) Emergent associative memory as a local organising principle for global adaptation in adaptive networks. *8th Int. Conf. on Complex Systems*. 417-430.
14. Mills, R.M., Watson, R.A. (submitted) Multi-scale search, modular variation, & adaptive neighbourhoods.
15. Okasha, S. (2006). *Evolution and the Levels of Selection*. Clarendon Press.
16. Parter, M. et al (2008) Facilitated Variation: How Evolution Learns from Past Environments to Generalize to New Environments. *PLoS Comput Biol* 4(11).
17. Penney, R.W., Coolen, A.C.C., Shenington, D. (1993) Coupled dynamics of fast spins and slow interactions in neural networks and spin systems. *J. Phys. A: Math. Gen.* 26:3681-3695.
18. Powers, S.T. (2010) *Social Niche Construction*, PhD thesis, ECS, Southampton.
19. Powers, S.T., Penn, A.S., Watson, R.A. (2011) The Concurrent Evolution of Cooperation and the Population Structures that Support it. *Evolution*. 65(6):1527–1543.
20. Sommer, R.J. (2009). The future of evo–devo: model systems and evolutionary theory. *Nature Reviews Genetics* 10 (6): 416–422
21. Szathmary, E. (2011) To Group or Not to Group? *Science.* 334 (6063) 1648-1649.
22. Toussaint, M., & von Seelen, W. (2007) Complex adaptation and system structure, *BioSystems* 90: 769–782
23. Watson, R.A. (2012) *Three Ways in Which the Algorithm of Biological Evolution Might not be Natural Selection (Even Though it is Based-on Natural Selection). ALife XIII*, to appear.
24. Watson, R.A., Buckley, C.L., Mills, R., and Davies, A. (2010) Associative Memory in Gene Regulation Networks. *Procs. 12th Int. Conf. on the Synthesis and Simulation of Living Systems (ALife XII)*.pp.194-202.
25. Watson, R.A., Buckley, C.L., Mills, R.M. (2010) Optimisation in 'Self-modelling' Complex Adaptive Systems. *Complexity.* 16(5):17-26.
26. Watson, R.A., Czapp, B., Powers, S.T., & Mills, R.M. (in prep.)*The Evolution of Ecological Relationships and Connectionist Principles of Adaptation in Ecosystems.*
27. Watson, R.A., Jackson, A., Palmius, N., Mills, R.M., Powers, S.T. (submitted) *The Evolution of Symbiotic Partnerships and Their Adaptive Consequences.*
28. Watson, R.A., Mills, R.M., Buckley, C.L. (2011) Global Adaptation in Networks of Selfish Components. *Artificial Life.* 17(3):147-66.
29. Watson, R.A., Mills, R.M., Buckley, C.L. (2011) Transformations in the scale of behavior and the global optimization of constraints in adaptive networks, *Adaptive Behavior* 19(4): 227-249.
30. Watson, R.A., Mills, R.M., Buckley, C.L., Powers, S.T., Penn, A.S., Davies, A., Noble, J, Bullock, S.G. (2010) Adaptation Without Natural Selection, in *Procs. ALife XII,* 80-81.
31. Watson, R.A., Palmius, N., Mills, R., Powers, S. and Penn, A. (2009) Can Selfish Symbioses Effect Higher-level Selection? *Procs. of 10th European Conference on Artificial Life (ECAL 2009).* 2:26-3