Emergent associative memory as a local organising principle for global adaptation in adaptive networks

Rob Mills, Richard A. Watson and C. L. Buckley
Natural Systems Research Group, University of Southampton
rob.mills@soton.ac.uk

Complex adaptive systems composed of self-interested agents can in some circumstances self-organise into structures that enhance global adaptation or efficiency. However, the general conditions for such an outcome are poorly understood. In contrast, sufficient conditions for artificial neural networks to form structures that perform collective computational processes such as associative memory/recall, generalisation and optimisation, are well-understood.

While such global functions within a single agent or organism may arise from mechanisms (*e.g.*, Hebbian learning) that were selected for this purpose, agents in a multi-agent system have no obvious reason to produce such global behaviours when acting from individual interest. However, Hebbian learning is actually a very simple and fully-distributed habituation or positive feedback principle. Here we use an adaptive network model in which agents can modify their behaviours (states) but also their interactions with other agents (network topology). We show that when self-interested agents can modify how they are affected by other agents then, in adapting these inter-agent relationships to maximise their own utility, they will necessarily alter them in a manner homologous with Hebbian learning. When the agents adapt their behaviours relatively quickly, and their relationships with other agents relatively slowly, we find that the overall network dynamics are modified to find better adapted states more reliably. This separation in timescales causes the state dynamics to spend most of their time at attractors. Thus, the network develops an associative memory that amplifies a subset of its own attractor states. This self-organised modification to the network dynamics enhances its ability to resolve conflicts between agents. Moreover, we show that the system is not merely 'recalling' high quality states that have been previously visited, but 'predicting' their location by generalising over local attractor states that have already been visited. Thus, globally adaptive behaviours can emerge from self-organising adaptive networks that follow organisational principles familiar in connectionist models of organismic learning.

1 Introduction

Many systems, both natural and artificial, are well described as complex networks. Understanding how complex networks composed of self-interested agents adapt and reorganise to exhibit globally efficient structures is in general an open question. Much research has been undertaken to consider this question in the study of systems that comprise adaptive components, and systems that have adaptive topology. More recently, interest has been increasing in systems in which both the component states and the network topology adapt, and indeed have reciprocal influence on one another [4]. They have been used to investigate the interplay between ecological dynamics and the slower-changing evolved inter-species relationships in ecosystems [13, 6], dynamics of social movement formation [3], and the evolution of cooperation [16, 14].

Hopfield networks, in the class of adaptive networks, are well-understood models of systems whose dynamics can be understood as the local minimisation of an energy function [7]. This energy minimisation ability can be interpreted as a optimisation process [8], in which the network finds locally optimal resolutions of the interactions between the system components (the strengths of these interactions can be used to encode an optimisation problem). However, due to the large number of local optima that arise from the presence of many constraints/interactions, the process will generally not result in a globally optimal resolution of those constraints [15]. A second and quite unrelated use of a dynamical system is to train the system to exhibit a particular energy function by changing the weights of interactions between system components, such that the dynamical attractors correspond to the configurations presented as training patterns. The system may then act as an associative or content-addressable memory [7]: when presented with a partial or noisy stimulus, the system 'recalls' the most similar training pattern.

Our recent work has brought together these seemingly disparate uses of dynamical systems, to show that organisational principles familiar in learing neural networks emerge spontaneously in these adaptive networks [18], and these locally mediated changes can lead to an enhanced ability to produce efficient system-level behaviours [17]. In this paper, we illustrate the mapping between evolving connections among selfish agents to model dynamics that arise from environmental pressures, and examine the sensitivity of the model to the separation in timescales of adaptation on the network and adaptation of the network topology.

2 A self-modelling adaptive network

We use a Hopfield network of N agents, each agent i accepting two behavioural states, $s_i = \pm 1$, the choice of which is determined by which state maximises individual pay-off, or utility, u_i:

$$u_i = \sum_{j}^{N} \omega_{ij} s_i s_j, \tag{1}$$

where $\omega_{ij} \in (-1, 1)$ is the strength of connection between agent i and agent j. States are updated asynchronously. Eqn 1 formally describes the function that an agent aims to maximise, which is achieved by minimising conflicts with the other agents. Specifically, agent i aims to align with all other agents j where $\omega_{ij} > 0$, and adopt the opposite state to agents j where $\omega_{ij} < 0$. Clearly in a given arbitrary system configuration not all of the conflicts can be simultaneously resolved. Thus, i will adopt the state that confers higher payoff given the current configuration of the other agents. We can describe the dynamics of the network by the updates to the individual states:

$$s_i(t+1) = \theta \left[\sum_{j}^{N} \omega_{ij} s_j(t) \right], \tag{2}$$

where θ is the Heavyside threshold function (which takes the value -1 for negative arguments and +1 for positive arguments). Hopfield showed that for symmetric weights and suitably constrained self-connections (here w_{ii}=1), the system dynamics only exhibit point attractors. These attractors are maxima in the system-level utility, U, that arises from the selfish actions of all the agents:

$$U = H(S(t), \Omega) = \sum_{i}^{N} u_i = \sum_{i} \sum_{j} \omega_{ij} s_i s_j, \tag{3}$$

where $S = (s_1, s_2, \ldots, s_N)$ is the vector of states of all the agents and Ω is a matrix of all the connection strengths. Thus, we can understand the dynamics of the network as a local optimisation process on the constraints of the network [8].

We consider the case of an adaptive network, in which the connection strengths are malleable. In some scenarios, agents may be able to directly modify their interactions with others (for example, in social networks [16]). In other cases, an agent may have influence over a set of connections that augment interactions that are external to the system. For example, epistasis between alleles is partly the result of external selective dependencies and partly the result of evolvable development. Thus, we suppose that ω_{ij} comprises the original, external dependencies, and the interconnections that are under an agent's control: $\omega_{ij} = \omega_{ij}^{E} + \omega_{ij}^{L}$.

As each agent chooses the state that maximises its own utility this may increase or decrease the utility of other agents in the network. We are interested in how this affects the total system utility of the system, which is simply the sum of individual utilities, measured using the environmental weights, ω_{ij}^{E}. Note that the state updates are a function of both the internal and environmental constraints, but when measuring the system-level behaviour we use just the environmental weights in order to enable direct comparison with non-adaptive controls.

To model the connection dynamics in a deterministic fashion, for a given agent i who has current internal weights ω_{ij}^{L} and is selected to undergo connection mutation, we produce two variant weights, $\omega_{ij}^{L+} = \omega_{ij}^{L} + \lambda$ and $\omega_{ij}^{L-} = \omega_{ij}^{L} - \lambda$, where λ is a constant that controls the rate of weight change. The weight that confers the highest payoff to agent i under Eqn. 1 is accepted (and replaces ω_{ij}^{L}).

The state of system takes on a random configuration every τ timesteps. In this paper, usually τ is sufficiently long to ensure that the system spends most of its time at an attractor, but sufficiently small compared to λ such that many attractors are visited in the time that significant modifications to connections takes place. There is no inheritance of system configuration after a perturbation, only the connection strengths are unmodified. Given sufficiently small λ, the cumulative effect of modifications to connections over a single relaxation of the network will be approximately equal to updating the connections (with a larger λ) only at the attractor. This end-of-relaxation update as used in our experiments is computationally less expensive to implement, but a continuous update model may be more natural in some systems.

Given Eqn. 1, it follows that when $u_i' > u_i$, and we have not allowed any state changes (thus $s_i' = s_i$ and $s_j' = s_j$), then [18]:

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

$$\Rightarrow \begin{cases} \omega_{ij}' > \omega_{ij} & \text{if } s_i s_j > 0, \text{ or} \\ \omega_{ij}' < \omega_{ij} & \text{if } s_i s_j < 0. \end{cases} \tag{4}$$

Intuitively, this equation tells us that for any pair of states that are aligned (*i.e.*, signs match), increasing the strength of the interaction increases individual payoff; while for any pair of states that are anti-aligned (*i.e.*, signs do not match), decreasing the strength of the interaction increases individual payoff. These directions of change are clearly homologous with Hebb's rule [5], which is often stated as $\Delta\omega_{ij} = \delta s_i s_j$ [1].

We assert a number of conditions on the operation of the network, which are sufficient to give rise to behaviours conferring global adaptation [17]:

1. that changes to both states and interaction weights are driven by selfish motives (of the components that are changing state or weight);
2. that the dynamics of the system have multiple point attractors;

3. that the system is repeatedly relaxed from different initial conditions, such that many local attractors are visited on the timescale where the interaction weights change slowly;
4. that the system dynamics are dominated by time spent at attractors.

The condition of self-interested changes (1) straightforwardly arises from the study of distributed complex systems. Without multiple attractors (2), there will be little to show in terms of *change* in the sizes of basins of attraction – for better or worse. The way in which we model the visitation of multiple attractors (3) here is by occasionally radically perturbing the states of all components. In §3.1 we illustrate the relationship between evolution in a fixed, multimodal environment and evolution to a set of unimodal environments that change over time (see also [12]). (4) ensures that changes to interactions largely occur when the network is at a local attractor (and hence agent's states are unchanging). This has the consequence that interaction changes that confer robustness (increase the stability of the current configuration) will be favoured. In §3.3 we investigate relaxing this assumption by reducing the separation in timescales.

2.1 Organisation of constraints in the environment

In this paper, we consider the environmental constraints under two different patterns of organisation:

- S1) highly organised, with strong modularity [11]. All within-module connections are of strength 1, and between-module connections are of strength p:

$$\omega^E_{ij} = \begin{cases} 1, & \text{if } \lfloor \frac{i}{k} \rfloor = \lfloor \frac{j}{k} \rfloor, \\ p, & \text{otherwise }, \end{cases} \tag{5}$$

- S2) random constraints, $\omega^E_{ij} \in [-1, 0, 1]$ nonzero with probability s, and ± with equal probability. There are Nk non-zero connections, so each agent has an average of k dependencies.

3 Simulated Experiments

In this section we provide investigations using our model with three distinct aims. The first part illustrates the mapping between training an artificial neural network and the memory that emerges in an evolving network of self-interested agents. The second part demonstrates how the self-interested actions of the agents leads to an increase in the ability of the network to resolve constraints between the agents. The third part examines the sensitivity of the model to the assumption that the timescales of agent state adaptation and connection adaptation are widely separated.

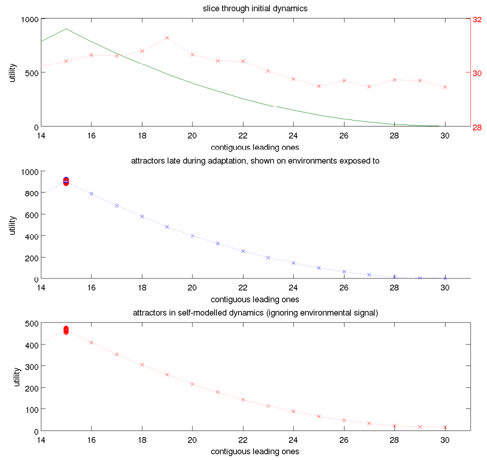

Figure 1: Slices through the attractor landscape and evolutionary landscape. Top) The initial dynamics are random (red dashes), with attractors unrelated to the environment (green). Middle) The attractors that the network visits during the last 25% of relaxations (red circles) are all at the environmental attractor. Bottom) When the environmental signal is removed, the network's internal dynamics recreate the same attractor (red circles). The red dashes show the utility of each configuration in this slice according to the internal dynamics only.

3.1 Selfish, local behaviours give rise to associative memory

Evolution in a unimodal environment

One basic ability that we would typically expect of a memory is to be able to increase the likelihood of a particular configuration – and in the context of a memory implicit in the dynamics of a network, that corresponds to an increase in the basin of attraction of that configuration. Here we allow the network to evolve under a simple unimodal environment, where the environmental weights are organised as per S1 (§2.1), with $N = 30, k = N/2, p = -0.9$. This creates a single attractor at $1^k(-1)^k$.[1] $\lambda = 0.01, \tau = 12N, R = 160$.

In Fig. 1, we show that when a network with random initial connections[2] is subjected to evolutionary pressures of this simple unimodal environment (top), the selfish changes to connection strengths leads to the network dynamics recreating a landscape whose optimum reflects the optimal configuration in the evo-

---

[1] Note the two-fold degeneracy of this form of spin-glass system means that $H(S) = H(-S)$, such that flipping all states is of identical utility; therefore we only plot the right-hand half of the utility surfaces

[2] drawn uniformly at random from $[-0.025, 0.025)$ and the connection matrix made symmetric

lutionary landscape (middle). Thus, even when we remove the evolutionary pressure (bottom), the internal network dynamics 'recall' the high-utility configuration in the environmental landscape.

Evolution with multiple attractors

As noted above, we are also interested in the ability to 'recall' multiple environments. We consider this in two different manners: first, with exposure to two unimodal environments in alternating episodes; and second, with a single multi-modal environment. The two environments are organised as per S1, with $N = 30, k = N/2$. The first environment, Ω^{E_1}, uses $p = -0.9$ with an attractor at $1^k(-1)^k$. The second environment, Ω^{E_2}, uses $p = 1$ with an attractor at 1^N. Each environment remains for $R_{\text{episode}} = 20$ relaxations. We set $\lambda = 0.05, \tau = 12N, R = 160$.

Fig. 2 (a) and (c) shows the network behaviour when subjected to these two alternating environments. As for the previous experiment, the initial random dynamics are unrelated to the environmental landscapes (frame a, top). Once connection adaptation has taken place, the network reaches the relevant attractor within most relaxations. However, occasionally the internal dynamics overpower the evolutionary pressure, and thus the evolutionary utility is very low (see frame a, middle, circle colour indicates which environment is current; see also (c), where a small number of circles during later episodes are of zero utility rather than that of the current attractor). When the evolutionary pressure is ignored, there are two attractors in the network dynamics, which correspond to the environmental targets (a, bottom).

A very similar result occurs under the bimodal environment to which the network is constantly exposed, $\Omega^{E_B} = 0.5\,(\Omega^{E_1} + \Omega^{E_2})$, which exhibits attractors in the same locations (see Fig. 2 (b) and (d)). Contrasting frames (c) and (d), we see that the distribution of attractors reached is more stochastic in (d), but the higher utility attractor of Ω^{E_2} has a slightly larger basin [2] and thus receives a larger proportion of visits. This higher visitation rate leads to positive reinforcement of the state, and over a longer period the higher-utility attractor will 'win out', becoming the only remaining attractor for the network. In additional experiments run over longer timescales (not shown), the higher-utility attractor does indeed become the only visited attractor. Conversely, when the mutation rate on connections is set too high, the emergent memory 'over-fits' to the attractor that is reached more frequently early in the trajectory, which may happen to be the lower utility attractor. Clearly such a case is undesirable, and thus for the emergent memory to have positive consequences the mutation rate on connections should not be too great.

3.2 Associative memory, when trained appropriately, can generalise to reveal better optima

In the previous two experiments, the total number of attractors is very small and consequently after very few relaxations all attractors are highly likely to

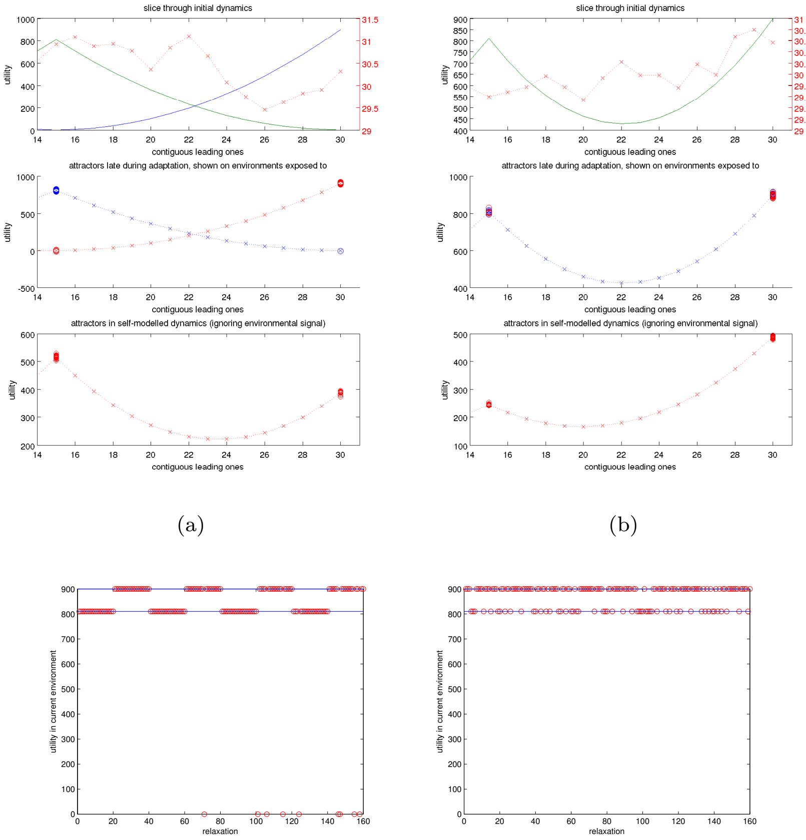

(a) (b)

(c) (d)

Figure 2: (a) Slices through the attractor landscape and the two unimodal evolutionary landscapes. (b) As per (a) for a single, bimodal evolutionary landscape that has attractors in the same locations.

be visited. In contrast, here we examine the changes to network dynamics in a multimodal environment with a very large number of attractors. We can assess the ability of the network to generalise and 'predict' novel combinations of previously seen sub-components, since the environmental landscape exhibits a greater level of ruggedness. Fig. 3 shows the distribution of the total utility at the end of each relaxation before and after evolution of the network connections, for 100 independent repeats on two environmental conditions: S2, with $N = 120, k = 8$; and S1, with $N = 100, k = 5, p = 0.01$.

This result shows that the selfish modification of connections results in enlargement of high-utility attractors at the expense of low-utility attractors, such that the total utility of equilibria found by the network with the modified connec-

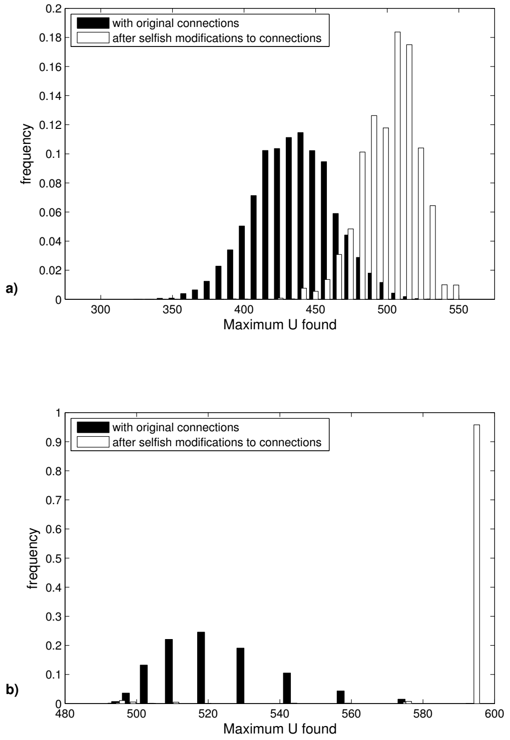

Figure 3: Total system utility measured before and after adaptation of connections, at the end of relaxation. a) a system of random interactions, 100 samples per run measured after 250 relaxations; and b) a system with strongly modular interactions between agents, 25 samples per run measured after 125 relaxations.

tions is higher on average than the total utility of equilibria found by the network with the original connections [18]. In the modular system the attractors of the system are known and we can estimate how the size of the basin of attraction for each attractor class (grouped by total utility) changes over time. There are many more low-utility attractors than high-utility attractors in this system and the global optimum is rarely visited initially. With unmodified connections, the time to first hit of the global optimum is 1009.8±92.00 relaxations (SEM) on average (over 100 independent experiments). In contrast, with slowly modified connections the time to first hit is 75.9±2.51 (SEM) (mean time to first hit of the global optimum is different with confidence p-value $< 10^{-5}$). This means that selfish modifications to connections are not merely increasing the basin of high-utility attractors that have already been visited, but are enlarging the basin of attraction for these attractors before they are visited for the first time.

3.3 On the importance of separating the timescales for adaptation *of* and *on* the network

As mentioned in §2, one of the conditions that we assume is that the network behaviour is dominated by steady state dynamics: that is, the states change relatively rapidly with respect to the weights/topology – but not *too* rapidly such that both processes are relevant to one another [6]. However, clearly not all systems exhibit such a separation in timescales, and accordingly the interaction may become more complicated [20]. Here we investigate the affordance of emergent associative memory when this assumption is relaxed.

Specifically, we vary the period between radical perturbations, τ, and identify a value of λ that results in the globally optimum being found from $\geq 95\%$ of initial conditions after selfish adaptation to connections. We use the highly organised environment S1, with $k = 5, N = 40$. Results are averaged over 30 independent repeats. The time it takes for the internal weights Ω^L to have a significant influence over the network dynamics is a function of $1/\lambda$, and accordingly, the lower the value of λ, the greater the number of relaxations we allow. In order to accurately measure the size of the basin of attraction of the global optimum, it is most straightforward to use an explicit measurement phase with long transients. That is, we evolve the network connections using short transients for some period, then disallow further changes to the network connections in order to measure the dynamics (using long transients of $\tau = 10N$).

Fig. 4 shows that for a large range of τ, we are able to find values of λ sufficient to result in high-quality global behaviour. At approx $\tau = 10N$, during adaptation is sufficient to observe reliable optimisation: the red and black curves converge (upper frame). Consider the extreme values of τ. With a very long separation in timescales, a local attractor will be reached early on in a trajectory, and thus the quality of information does not increase even though the computing time does increase. With very short transients, the signal of what configurations are compatible is very noisy, and accordingly the number of samples that is required to identify the signal amongst the noise is very large. However, the minimal time spent adapting (bottom panel) is on the order of $\tau = N$. This tell us that, for these environmental parameters, there exists a favourable trade-off in the cost of moving to a local attractor and the 'quality of information'.

Of course, as a model of a biological system, what is important is that the emergent associative memory has a positive effect on the global dynamics for a broad range of τ values. Previously, we could reason that any trajectory of length sufficient to reach a local attractor with reasonable likelihood, denoted as τ^*, would confer a positive effect. This result shows that there exist environmental scenarios in which the separation in timescales can be $\tau < \tau^*$ and the associative memory that emerges will have a positive effect in terms of global constraint resolution.

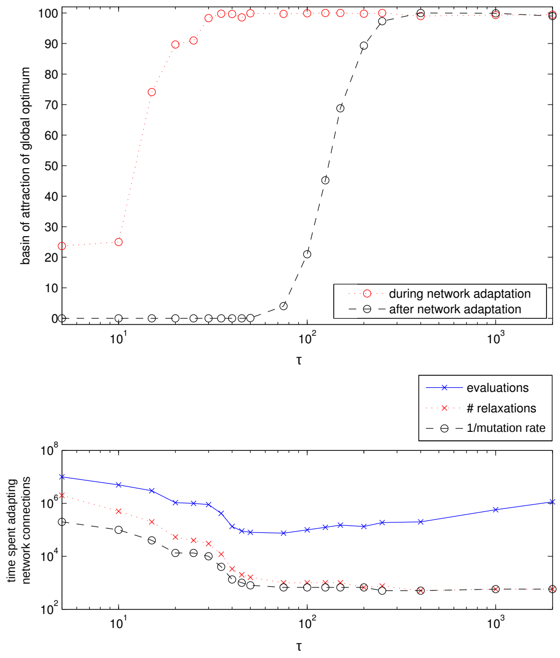

Figure 4: System performance as a function of the separation in timescales. Upper: size of basin of attraction of the highest-utility attractor after reorganisation. Lower: mutation rate on connections used to achieve this reorganisation.

4 Discussion

The results above provide an illustration of the mechanisms and capabilities of an associative memory that emerges from the action of self-interested agents in an adaptive network. Where an artificial learning neural network may be explicitly presented with patterns from which to train, in our model the analogous patterns come from the system's own dynamics (which themselves are a product of both the environmental pressures and the internal constraints). The recall behaviour of an artificial neural network is typically invoked with the presentation of a partial stimulus, and the analog in our model is to run the dynamics from a new initial condition, resulting from some form of environmental change. Above we illustrated the equivalence between exposure of the network to a set of simple environments and a constant but rugged utility landscape. In the latter formalism, we encode a complex scenario in the environmental weight matrix, and we find it appropriate to (partially) randomise the states periodically. The presentation of varying environments [12] over time has a qualitatively similar result, and may also implicitly provide the occasional radical perturbations as explicitly modelled above. We also have some preliminary evidence to suggest that the scale of the perturbations need not be as radical as to randomise all of the agent's states.

In this paper we investigated a simple and distributed method for forming an associative memory, examining the consequences with reference to the ability of the network to lead to high-utility states at the global scale. Elsewhere, we have developed a similar model that uses the same protocol to develop associations, but uses the evolved interaction weights in a different manner. Specifically, where here we interpret the evolved interactions as a modification on the selection operation, elsewhere we consider how the system behaviour can be transformed when the evolved interactions inform correlated variational steps in the dynamics. Thus, in the alternative work, associations between components correspond to the likelihood of co-variation [19, 9]. We have shown that this alternative method of using the information can lead to a transformation in the scale of dynamics – effectively moving to a lower-dimensional version of the original system. We have examined the algorithmic affordances of this method of team formation [10], demonstrating rapid problem-solving ability in problems with multi-scale structure.

A number of aspects of our model warrant further investigation. For instance, two assumptions that are made in the current work are: 1) that the periodic perturbations or shocks are radical, *i.e.*, they randomise the state on every component in the system; and 2) that the initial constraints on the system are symmetric. The latter assumption ensures that the dynamics only exhibit point attractors, whereas asymmetric constraints can lead to limit cycles or chaotic attractors. While we do not believe that the strictness of either of these assumptions is required for our result, the degree to which each can be relaxed remains to be investigated.

In conclusion, we have shown that simple self-interested adaptation of agents can lead to the emergence of an associative memory, when those adaptations occur on two levels – states and connections/network topology. When these two processes occur on timescales that are sufficiently separated, the emergent memory effectively facilitates the capability to predict the position of higher quality attractors than are discovered by a system that lacks the ability to modify its own topology. In this work, we have also shown that our previous idealisation that adaptation of connections should occur at or close to local attractors in the system dynamics can be relaxed. This indicates that dynamical transients that have moved some way towards attractors provide information that result in agent's selfish decisions being beneficial for efficiency at the global scale, 'merely' as a side effect. The benefits of an emergent associative memory can be seen at the system level through the ability to generalise over the state configurations that the system has experienced, and in some circumstances 'predict' the location of high-quality attractors that were previously unvisited.

Bibliography

[1] Ackley, D. H., G. E. Hinton, and T. J. Sejnowski, "A learning algorithm for boltzmann machines", *Cognitive Science* **9** (1985), 147–169.

[2] Coolen, Ton, "On the relation between stability parameters and sizes of domains of attraction in attractor neural networks.", *Europhys Lett* **16** (1991), 73.

[3] Geard, Nicholas, and Seth Bullock, "Competition and the dynamics of group affiliation", *Advances in Complex Systems* **13**, 4 (2010), 501–517.

[4] Gross, Thilo, and Bernd Blasius, "Adaptive coevolutionary networks: a review", *Journal of The Royal Society Interface* **5**, 20 (2008), 259–271.

[5] Hebb, Donald O., *The organization of behavior: a neuropsychological theory*, Wiley New York (1949).

[6] Hogeweg, Paulien, "From population dynamics to ecoinformatics: Ecosystems as multilevel information processing systems", *Ecological Informatics* **2**, 2 (2007), 103–111.

[7] Hopfield, John J., "Neural networks and physical systems with emergent collective computational abilities", *PNAS* **79** (1982), 2554–2558.

[8] Hopfield, John J., and D. Tank, "Computing with neural circuits: A model", *Science* **233** (1986), 625–623.

[9] Mills, Rob, *How Micro-Evolution Can Guide Macro-Evolution: Multi-Scale Search via Evolved Modular Variation*, PhD thesis School of Electronics and Computer Science, University of Southampton (2010).

[10] Mills, Rob, and Richard A. Watson, "Multi-scale search, modular variation, and adaptive neighbourhoods", *(under review)*.

[11] Mills, Rob, and Richard A. Watson, "Variable discrimination of crossover versus mutation using parameterized modular structure", *Proceedings of the Genetic and Evolutionary Computation Conference* (H. L. et al ed.), ACM Press (2007), 1312–1319.

[12] Parter, M, Nadav Kashtan, and Uri Alon, "Facilitated variation: How evolution learns from past environments to generalize to new environments", *PLoS Computational Biology* **4**, 11 (2008), e1000206.

[13] Poderoso, Fábio C., and José F. Fontanari, "Model ecosystem with variable interspecies interactions", *Journal of Physics A* **40** (2007), 8723–8738.

[14] Santos, Francisco C, Jorge M Pacheco, and Tom Lenaerts, "Cooperation prevails when individuals adjust their social ties", *PLoS Computational Biology* **2** (2006), e140.

[15] Tsirukis, A.G., G.V. Reklaitis, and M.F. Tenorio, "Nonlinear optimization using generalized hopfield networks.", *Neural Computation* **1** (1989), 511–521.

[16] Van Segbroeck, Sven, Francisco Santos, Ann Nowe, Jorge Pacheco, and Tom Lenaerts, “The evolution of prompt reaction to adverse ties”, *BMC Evolutionary Biology* **8**, 1 (2008), 287.

[17] Watson, Richard A., C.L. Buckley, and Rob Mills, “Optimisation in ‘self-modelling’ complex adaptive systems”, *Complexity* (2010), (in press).

[18] Watson, Richard A., Rob Mills, and C.L. Buckley, “Global efficiency in networks of selfish components: emergent associative memory at the system scale”, *Artificial Life* (accepted 2011).

[19] Watson, Richard A., Rob Mills, and C.L. Buckley, “Transformations in the scale of dynamical behaviour and the global optimisation of constraints in adaptive networks”, *Adaptive Behavior* (accepted 2011).

[20] zu Erbach-Schoenberg, Elisabeth, Connor McCabe, and Seth Bullock, “On the interaction of adaptive timescales on networks”, *Procs. ECAL*, (2011 (under review)).