Associative memory in gene regulation networks

Richard A. Watson[1], C.L. Buckley[2], Rob Mills[1], Adam Davies[1]

[1]Natural Systems group, ECS, University of Southampton, U.K. [2]CCNR, Informatics, University of Sussex, U.K..
raw@ecs.soton.ac.uk

Abstract

The pattern of gene expression in the phenotype of an organism is determined in part by the dynamical attractors of the organism's gene regulation network. Changes to the connections in this network over evolutionary time alter the adult gene expression pattern and hence the fitness of the organism. However, the evolution of structure in gene expression networks (potentially reflecting past selective environments) and its affordances and limitations with respect to enhancing evolvability is poorly understood in general. In this paper we model the evolution of a gene regulation network in a controlled scenario. We show that selected changes to connections in the regulation network make the currently selected gene expression pattern more robust to environmental variation. Moreover, such changes to connections are necessarily 'Hebbian' – 'genes that fire together wire together' – i.e. genes whose expression is selected for in the same selective environments become co-regulated. Accordingly, in a manner formally equivalent to well-understood learning behaviour in artificial neural networks, a gene expression network will therefore develop a generalised associative memory of past selected phenotypes. This theoretical framework helps us to better understand the relationship between homeostasis and evolvability (i.e. selection to reduce variability facilitates structured variability), and shows that, in principle, a gene regulation network has the potential to develop 'recall' capabilities normally reserved for cognitive systems.

Evolvability

How natural selection results in the evolution of complexity, if it is natural selection that is responsible, is not yet understood [1,2]. It is easy to see how natural selection increases the frequency of fit phenotypes from a given distribution of phenotypic variants. But this is only part of the explanation. Although continued adaptation does not require that the available distribution of phenotypes is fitter than the parent on average (that would imply directed variation), continued increases in fitness and functionality require that this distribution includes at least some phenotypes that are fitter than the parent. This is often taken for granted, but experience in evolutionary algorithms and artificial life experiments suggests that such variants are quickly exhausted by selection, precluding further adaptation [2]. Thus the evolution of significant biological complexity requires that we explain how the distribution of phenotypes, resulting as they do from random variation in genotypes, includes phenotypes that are, not merely different from, but fitter than the parental type. The explanation might be, at least in part, that in natural organisms the distribution of phenotypic variants itself becomes better adapted over time [3] – hence enhancing *evolvability*, the ability of a population to evolve [4,5,6,7]. Since the processes of development, mapping genotype to phenotype, is itself genetically specified and subject to natural selection, this seems like a possibility, at least in principle.

However, although it is easy to say that natural selection should favour more evolvable genotypes, without a proximal account for the selective gradients that would produce such an outcome this is just wishful thinking. It is not so easy to pin down the source of a selection pressure that increases evolvability. For example, enhanced evolvability ought to mean that a genotype evolves *better*, not just that it evolves, and given that adaptive variants from a given phenotypic distribution are quickly exhausted it is hard to see how a variant genotype in a population that is stuck at a local optimum can be said to have better evolvability than another. This implies that the evolution of evolvability might require a constantly varying selective environment and multiple opportunities to generate and exploit variant phenotypic distributions. Moreover, if the environment changes in an entirely arbitrary fashion, a genotype to phenotype mapping cannot evolve to exploit it, so we are lead to the conclusion that such a mapping could only be adaptive if it exploits some kind of structure or regularity observed in the distribution of selective environment [8].

A simple way in which this might work is as follows. Different genotypes with the same phenotype might (nonetheless) have a different distribution of *phenotypic neighbours* - phenotypes produced through small mutations to the genotype. In a selective environment that varies from one selective regime to another (Fig.1), natural selection might favour genotypes that have phenotypes that are fit in one regime *and* have phenotypic neighbours that are fit in the other (over genotypes that have phenotypes that are equally fit in the first regime but do not have phenotypic neighbours that are fit in the other) [8]. In a sense, we can understand the propensity to produce phenotypes that are not currently selected for but have been selected for in the past as a kind of 'memory' of past selective environments [8], and under certain conditions evolved genotypes may even "generalise to future environments, exhibiting high adaptability to novel

goals". But exactly how this might happen, what the selective pressures are that might produce this outcome, and the limitations and affordances of such a process are poorly understood in general.

Part of the process might involve the evolution of modularity, for example [9,10]. That is, certain phenotypic features might become tightly integrated units (clusters of phenotypic features that co-vary), whilst others remain, or become, separated and vary independently. Such modularity might then provide, in effect, higher-level variation – i.e. variation at a higher-level of organisation [11]. Such high-level variability might in principle provide new combinations of modules with high probability (compared to the original distribution of 'atomic' character combinations) even though some particular combination of modules that is fit may not previously have been selected for.

Wagner *et al* [10] explain part of the proximal mechanism that might be involved in this process. Referring to genetic loci that affect the correlation of phenotypic traits [12], they state that "natural selection can act on [such loci] to either increase the correlation among traits or decrease it depending on whether the traits are simultaneously under directional selection or not. …[Resulting in] a reinforcement of pleiotropic effects among co-selected traits and suppression of pleiotropic effects that are not selected together" [10].

Wagner *et al* do not seem to notice, however, that this suggests intriguing parallels with Hebbian learning familiar in computational neuroscience [13,14]. Hebb's rule, in the context of neural network learning, is often represented by the slogan *neurons that fire together wire together*, meaning that synaptic connections are strengthened between neurons that have correlated activation in response to a stimulus. Formally, a common simplified form of Hebb's rule states that the change in a synaptic connection strength ω_{ij} is $\Delta\omega_{ij} = \delta s_i s_j$ where $\delta>0$ is a fixed parameter controlling the learning rate and s_n is the current activation of the n^{th} neuron. This learning rule has the effect of transforming correlated neural activations (created by an external stimulus) into causally linked neural activations. From a dynamical systems perspective, this has the effect of enlarging the basin of attraction for the current activation pattern/system configuration created by the stimulus. This type of learning can be used to train a recurrent neural network to store a given set of training patterns [15] thus forming what is known as an 'associative memory' of these patterns. A network trained with an associative memory then has the ability to 'recall' the previously seen training pattern that is most similar to a new partially specified or corrupted test pattern.

In this paper we investigate the possibility that a gene regulation network, capable in principle of exhibiting the same kind of dynamics as a recurrent neural network, is subject, over evolutionary timescales (not lifetimes [16]), to modifications in connections that are in principle the same as those produced by Hebbian learning familiar in neural network models. Thus *genes that fire together wire together* - i.e. genes whose expression is selected for in the same selective environments become co-regulated. Accordingly, the previously external cause of correlations in phenotypic characters (i.e. direct selection on expression patterns) becomes internalised (i.e. the result of a regulatory connection). A developmental trajectory determined by such an evolved network will then be able to reproduce a previously selected phenotype ballistically from an arbitrary initial condition using purely internalised dynamics, i.e. using a memory of what phenotypic characters work well together.

This analogy helps us to understand how a gene regulation network can modify the distribution of phenotypes in a manner that reflects structure in the selective environment. Specifically, we argue that evolved changes in regulatory connections will tend to cause the regulatory network as a whole to form an associative memory [15] of locally optimal phenotypes that have been visited in the past [17,18]. The evolved network has a dynamical behaviour which models the historical selective pressures on phenotypes (in the sense of having the same attractors) and can thereby create phenotypic distributions that are especially fit. In particular, an evolved network can produce a distribution of phenotypes that enables a population to escape locally optimal phenotypes (i.e. phenotypes that were locally optimal prior to the development of this regulation) in favour of superior optima. We also show that the proximal cause of these changes is not the teleological anticipation of future reward but something much more mundane – merely selection for robustness or canalisation of the current phenotype [5]. By analogy with the Baldwin effect [19], the internalised memory of previously found solutions enables previously evolved phenotypes to be produced innately by the developmental process. We therefore argue that selection for homeostasis on an immediate timescale (i.e. the ability to regulate a constant condition [20]), is the proximal cause of increased evolvability on larger timescales (i.e. increased ability for adaptation), as we will discuss.

Self-modelling dynamical systems

In related work [17,18] we have been developing the concept of a 'self-modelling' dynamical system – a complex adaptive system that creates a memory of its past dynamical behaviour. We have shown that if changes to connections are Hebbian and slow compared to the system's state dynamics, a complex adaptive system will form an associative memory of its own dynamical attractors that enables it to lower its energy more efficiently and completely when subjected to repeated perturbation [17]. The 'training patterns' in such a scenario are the configuration patterns that are commonly experienced under the network's intrinsic dynamics, hence '*self*-modelling' [18] – and if the system spends most of its time at locally optimal configurations, it is these configurations that the associative memory stores. From a neural network learning point of view, a network that forms a memory of its own attractors is a peculiar idea. Forming an associative memory means that a system forms attractors that represent particular patterns or state configurations. For a network to form an associative memory of its own attractors therefore seems redundant; it will be forming attractors that represent attractors that it already has. However, in forming an associative memory of its own attractors the system will nonetheless alter its attractors; it does not alter their positions

in state configuration space, but it does alter the size of their basins of attraction (i.e. the set of initial conditions that lead to a given attractor state via local energy minimisation). Specifically, the more often a particular state configuration is visited the more its basin of attraction will be enlarged and the more it will be visited in future, and so on. Because every initial condition is in exactly one basin of attraction it must be the case that some attractor basins are enlarged at the expense of others. Accordingly, attractors that have initially large basins of attraction will, with continued positive feedback, eventually out-compete all others until there is only one attractor remaining in the system.

Variation in the selective targets/initial conditions

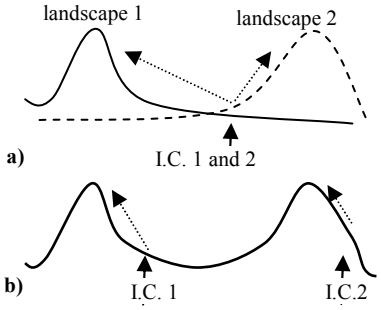

Fig.1. a) Adaptation to two different targets from the same initial condition (I.C.), b) Adaptation to one multi-modal target from two different initial conditions.

Before introducing our model, we briefly discuss an equivalence between multiple evolutionary episodes in different selective environments (Fig.1.a) and multiple evolutionary episodes from different initial conditions in a static (but multi-modal) selective environment (Fig.1.b). Parter *et al*, for example, conduct experiments using the former – and construct by hand different selective targets that are drawn from the same 'language' of tasks [8] (varying in a modular manner). We prefer the latter; using a single multi-modal landscape (created by modular epistasis) with repeated radical 'perturbations' of the evolved solution causing it to visit different local optima. What matters for our purposes is only the similarity or differences of the multiple 'targets'/ 'local optima', and the latter method has the advantage that, when the landscape is produced from the superposition of many low-order epistatic interactions (see methods), it does not require such explicit hand-crafting in this respect since structural similarity in the local optima results naturally.

A model for the concurrent evolution of gene expression patterns and regulation networks

Overview. Our model is intended to be as simple as possible. Presumably, the evolution of a gene expression network that is capable of creating correlated gene expression patterns and potentially sophisticated dynamical attractors was preceded by the evolution of static (unregulated) gene expression patterns. Likewise, the evolution of robust cell types in single-celled organisms, and gene expression networks that (partially) determine those cell types, presumably preceded the evolution of multi-cellular development and programmed cell differentiation. Accordingly, our model addresses the evolution of a gene expression pattern, and subsequently a regulation network, in a single-celled organism. By 'phenotype' we therefore simply mean a particular pattern gene expression, and by 'development' we simply mean the dynamical gene regulation process that creates the 'adult' gene expression pattern.

The model is not intended to be a literal model of biological processes. The critical features include a continuous-valued state vector representing a pattern of gene expression and a matrix of positive and negative connections representing up- and down-regulating connections between genes. These are subject to random variation and a selective environment that favours particular gene expression correlations. These components are linked together in a manner representing the concurrent evolution of a gene expression pattern and a gene regulation network but we aim to keep this protocol as simple as possible (see Fig. 2).

We assume that a pattern of gene expression is (epigenetically) inherited from one cell to the descendant cell and that a selection pressure on this phenotype causes it to evolve over many reproductions. A regulation network is also (genetically) inherited and subject to evolution via selection on the gene expression pattern that it modifies. We assume that every gene has the potential to regulate any other gene but that there is no significant regulation in the ancestral cell type (i.e. initially zero connections). Random variation in the connections of the network can introduce positive or negative correlations in the expression of genes which may or may not be beneficial given the current selective environment. So, in the lifetime of the cell, its initial gene expression pattern is inherited from the parent cell with random variation, this pattern of expression then forms the initial condition of the gene regulation network, which is then run for a number of time-steps (usually one) creating a slightly altered pattern of gene expression, and it is this pattern of expression which is interpreted as the phenotype of the organism and evaluated by the fitness function.

Evolutionary adaptation. The idea of evolved *correlations* between the expression of one gene and that of another invokes the notion of a *distribution* of phenotypes. When there are many copies of each genotype in a population, each one producing a phenotype from this distribution, selection on these individual phenotypes implicitly selects for genotypes that produce high fitness phenotype distributions [10]. However, we find that an explicit population with multiple copies of a genotype is more complicated than necessary. It is sufficient to merely compare the phenotype of a mutant to the phenotype of the original type and retain whichever is fitter. Hence we model the evolutionary process with a simple random mutation hill-climber (or '(1+1)ES'[21]) rather than a population-based evolutionary algorithm [3]. The latter merely adds additional stochastic fluctuations and unnecessary conceptual complications.

The overall architecture of the evolutionary model is depicted in Fig. 2. and detailed in Fig.3. Note that the gene

expression network does not so much represent a mapping from genotype to phenotype, as it is popularly conceived, so much as a mapping from an initial gene expression pattern to an 'adult' gene expression pattern. This adult gene expression pattern and the gene expression network is passed on the next generation (with random variation).

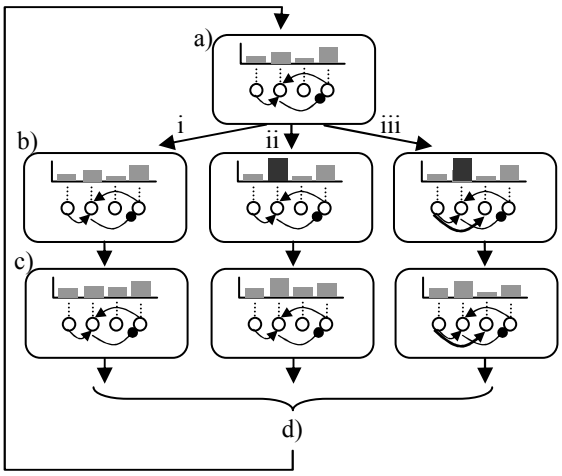

Fig.2: Schematic overview of the inheritance, regulation and selection processes (i.e. an iteration of the evolutionary hill-climber). a) A cell contains both an expression pattern and a genetically specified gene regulation network. b) Its descendents include individuals that are i) identical to the parent, ii) have a perturbed expression pattern (black), iii) have both a perturbed expression pattern and a genetically mutated regulation network (here depicted by an additional connection). c) The pattern of gene expression in each of these descendent cells is 'developed' or 'run' through their regulation networks creating three slightly different 'adult' gene expression patterns. d) The cell with the most fit gene expression pattern replaces the ancestral cell type.

The gene regulation network, *R*, (Fig. 3) is a matrix of connection strengths initialised to 0. The expression pattern, *E*, is set to a random configuration each t^*=5000 iterations (each gene expression level is set to a value drawn uniformly and independently in the range (-1,1)). This represents a radical environmental perturbation of the expression pattern and allows the expression pattern to visit the slopes of different local optima in the fitness landscape (Fig. 1) hence commencing a new evolutionary 'episode'. *E1*, *E2* and *E3* are the three modified expression patterns that result from the three descendents of the ancestral type (having no mutations, mutation to the expression pattern only, and mutation to both the expression pattern and the regulation network, respectively. We assume that mutation to the regulation network without mutation to the regulation pattern is unlikely). *mut* is a mutation function that introduces a small perturbation to the expression pattern or a small mutation to the regulation network. Specifically one of the existing expression levels or connection strengths (selected at random) is modified by adding a value drawn uniformly in the range (-1,1). (In test cases where the regulation network is not evolved, lines 2.c and 2.g are omitted.) *run*(*E,R*) is a function that 'develops' the initial expression pattern *E* by running the regulation network *R* for *p* time steps (*p*=1 by default) and returns a new expression pattern. For each time step the new activation level, s_i(t+1), of gene, *i*, is calculated using the old value with a decay term and a sum of weighted (positive or negative) inputs from the other genes in the network, as follows [22]:

$$s_i(t+1) = s_i(t) + T\left(\sum_j^N w_{ij}\sigma(s_j(t)) - s_i(t)\right) \quad \textbf{(1)}$$

where T=0.001 is a time constant, w_{ij} is the connection from gene *j* to gene *i*, $\sigma(x)=\tanh(x/10)$ is a sigmoidal output function determining the expression level of a gene with activation level *x* (representing the tendency of expression levels to saturate).

1. initialise regulation network, *R*.
2. *t*=0, repeat
   a. if (*t*=0) expression pattern, *E*=random, *t*=*t**;
   b. E'=*mut*(*E*);
   c. R'=*mut*(*R*);
   d. E1=*run*(*E*, *R*); E2=*run*(*E*', *R*); E3= *run*(*E*',*R*')
   e. *m*= max(*f*(*E1*),*f*(*E2*),*f*(*E3*))
   f. if (*f*(*E2*)=*m*) *E*=*E*';
   g. if (*f*(*E3*)=*m*) *E*=*E*', *R*= *R*';
   h. *t*=*t*-1

Fig. 3. Pseudocode of the inheritance, regulation and selection processes depicted in Fig. 2.

The selective environment. The fitness landscape is (initially) carefully controlled so that we can assess easily whether an evolved regulation network is creating appropriate correlations in the gene expression pattern. The minimal conceivable scenario is one where there are only two genes with selection for correlated expression in these two genes [10]. If we do not have any intrinsic preference for absolute gene expression levels, only for correlations, this means that there will be two locally optimal gene expression patterns of equal fitness – 'HH' and 'LL' (representing 'High' or 'Low' expression levels for the first and second genes). Alternatively, if we select for anti-correlation then these will be 'HL' and 'LH'. However, although we might be able to evolve a gene regulation network that supports correlation or anti-correlation in such a scenario, the evolutionary outcome will be somewhat degenerate in the sense that each of the two locally optimal gene expression patterns will have equal fitness and be equally likely to arise (from a random initial condition) without a regulation network.

Accordingly, we will examine the next simplest case; a system of four genes in two pairs. Here we can define a fitness function where 'HHHH' and 'LLLL' are maximally fit, but where 'HHLL' and 'LLHH' are local optima of lower fitness. Favouring pairs of co-expressed genes in this manner thus enables us to define a system with different-fitness optima without introducing a preference for absolute expression levels, or any asymmetries that would make one gene more important than any other. It also represents a minimally 'modular' fitness function. Naturally, we do not imagine that such a fitness landscape represents any realistic biological scenario – its structure is chosen merely to avoid obfuscating the significance of an evolved regulation network

with a complex adaptive landscape, and to test whether a network can create correlations that support co-regulation and create high-fitness phenotypes (we later investigate evolution on a 30-variable randomised landscape).

We construct a fitness function of this type using a sum of low-order (pair-wise) epistatic interactions [23] creating a locally smooth (but multi-modal) fitness landscape. Specifically, the fitness of an expression pattern, $S=<s_1,s_2\ldots s_N>$, is given by:

$$f(S) = \sum_{i}^{N}\sum_{j}^{N} e_{ij}\sigma(s_i)\sigma(s_j) \qquad \textbf{(2)}$$

where N is the number of genes in the system, s_i is the activation of the i^{th} gene, e_{ij} is the epistatic interaction between genes i and j, defined below and $\sigma(s)=\tanh(s/10)$ is the expression level of the gene, as before. The epistatic matrix is as follows: $e_{12}=e_{34}=1$, $e_{13}=e_{14}=e_{23}=e_{24}=0.1$, else $e_{ij}=0$ – thus defining the two pairs of strongly interacting genes (s_1/s_2 and s_3/s_4), with only weak interactions between these pairs as discussed above.

Results

Evolution of expression patterns without evolved regulation. Fig. 4 (right) illustrates the evolution of an expression pattern (without evolved regulation) over 10^5 evolutionary time steps (therefore showing 20 evolutionary episodes between radical perturbations of the expression pattern). This clearly shows the four locally optimal expression patterns (HHHH, HHLL, LLHH, and LLLL) and that patterns where the four genes are all high or all low have the highest fitness. The fitness values at each of the evolutionary local maxima attained (i.e. at each t=1 time step) may be either in the lower class or the higher class (see Fig. 4). The proportion of high and low fitness optima found indicates the size of the evolutionary basin of attraction for each class of optima. For these parameters under these conditions (without a regulation network) we find that the evolutionary basin of attraction for the fitter local optima accounts for about 73% of the initial configuration space (averaged over 300 evolutionary episodes).

Evolved regulation. Under natural selection, evolved changes to the connections in the regulation network must be those that change the expression pattern in the direction that increases fitness; and that direction may be different depending on the currently selected expression pattern. Since the evolved expression pattern very quickly settles into one attractor or the other, most evolution of the regulation network will occur when the expression pattern is at or near a locally optimal configuration. So, as a first step to investigating the evolution of a regulation network we evolve the regulation network when the expression pattern is 'clamped' at a single locally optimal configuration. Specifically, in line 2.a of Fig.3, E is set to $<s,s,s,s>$ (s=5) instead of a random configuration. We find that after 100,000 more evolutionary steps the evolved connections in the regulation network are all positive (Table 1). In contrast, when the clamped expression pattern is HHLL ($E=<s,s,-s,-s>$), the evolved connections are positive on the block diagonal (shaded) and negative elsewhere (Table 2).

It is crucial to note that the signs of these connections do not directly reflect the epistatic interactions in the fitness landscape – the intrinsic epistasis in the landscape does not change between the HHHH and HHLL test cases. Rather the evolved connections reflect the expression states experienced when the regulatory connection is altered (i.e. s_i=H/s_j=H and s_i=L/s_j=L expression levels create selection for positive connections, whereas s_i=H/s_j=L and s_i=L/s_j=H expression levels evolve negative connections). This clearly follows Hebbian principles – when equal gene expression levels are selected together they wire together positively, when one is selected to be high and the other low, they wire together negatively.

However, the sign of the connection is really just a labelling convention – what really matters with respect to demonstrating Hebbian learning is that these evolved connections increase the basin of attraction for the current expression pattern. Fig. 5 shows, for example, the effect of the connections evolved at the HHLL expression pattern (i.e. Table 2). We see that the evolved connections change the size of the HHLL attractor basin to fill 100% of the configuration space (conversely, when regulation is evolved at the HHHH expression pattern, Table 1, this pattern comes to occupy 100% of the configuration space).

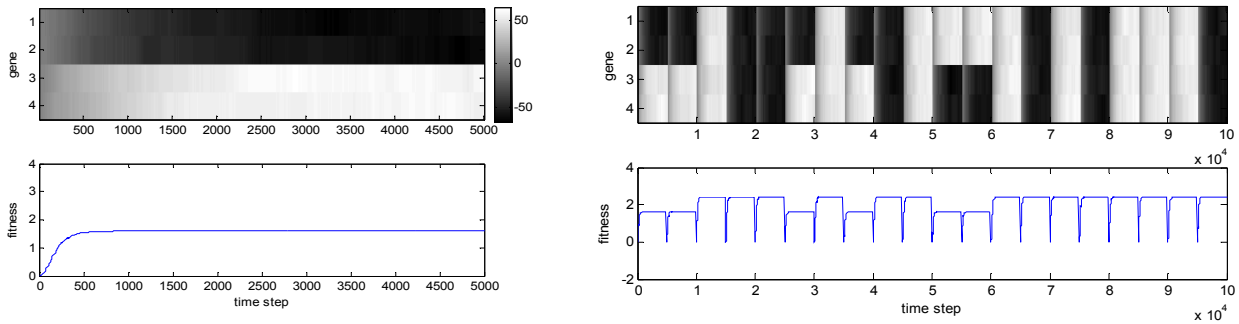

Fig.4. left) Evolution of a gene expression pattern without regulation for one evolutionary episode (5000 time steps). This happens to arrive at the locally optimal expression pattern where genes 1 & 2 are low, and 3 & 4 are high. Right) A longer run (100,000 time steps) including 20 evolutionary episodes, again without evolved regulation. Note that with these parameters, each evolutionary episode very quickly reaches a locally optimal expression pattern (i.e. transients are short). Note that fitnesses at evolutionary attractors fall into two classes (roughly those below a fitness of 2 and those above).

i/j	1=H	2=H	3=H	4=H
1=H	89.13	160.18	126.02	104.35
2=H	120.42	58.95	87.40	152.94
3=H	163.49	76.60	152.08	79.10
4=H	197.69	56.58	158.36	159.87

Table 1: evolved connections when the expression pattern is HHHH.

i/j	1=H	2=H	3=L	4=L
1=H	80.93	105.81	-60.99	-146.92
2=H	153.02	120.27	-94.84	-108.03
3=L	-157.65	-125.27	69.33	163.97
4=L	-156.00	-140.19	84.13	69.17

Table 2: evolved connections when the expression pattern is HHLL.

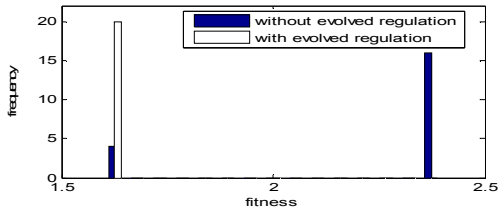

Fig. 5. Number of evolutionary episodes (from 20) finding each locally optimal phenotype before and after evolution of the regulation network. When the gene expression pattern is held at a low fitness attractor, the evolved regulation network canalises this pattern.

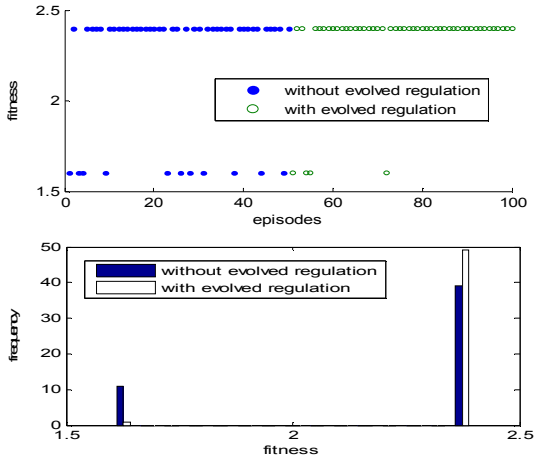

Fig. 6. When the gene expression pattern is evolved freely, evolved regulation canalises the fitter pattern (since it is visited more often). Upper) The evolution of a gene expression pattern without evolvable regulation (episodes 1-50) and with evolvable regulation (episodes 51-100). Each point represents a locally optimal expression pattern found via a single evolutionary episode from a random initial condition. Lower) see Fig.5.

i/j	1	2	3	4
1	437.37	566.40	60.50	72.32
2	269.72	389.88	253.21	212.56
3	184.52	98.54	270.58	351.04
4	448.46	-25.23	373.18	246.46

Table 3: Evolved regulatory connections when the expression pattern is not clamped. Although there is a lot of variation, the average value in the block diagonal (shaded) is 363 and elsewhere 163. The generally positive values mean that both the superior HHHH/LLLL attractor (Table 1) and the inferior HHLL/LLHH attractor (Table 2) have been reinforced, but the lower values off the diagonal retain a reflection of the underlying modularity.

Note that the evolved regulation network does not necessarily increase the basin of attraction for the fitter phenotypes, but rather for the phenotype present at the time that changes to the regulation network were evolved. Next, we evolve the regulation network without clamping the expression pattern. Without regulation the fitter phenotype is already found 73% of the time, so if the evolved regulation network reinforces the fitter attractor 73% of the time and the less fit attractor only 27% of the time then on average the fitter attractor should be enlarged more often than the less fit attractor in a positive feedback manner and it will eventually outcompete it (Fig. 6, Table 3).

Collectively, these results demonstrate that selection favours changes to regulation connections that reflect co-expression in the current phenotype, and that these connections increase the basin of attraction for that expression pattern, as expected for Hebbian changes to connections. They also show that in a fitness landscape where fitter patterns have larger basins (as is necessarily the case when the fitness landscape is created from the superposition of many low order interactions [18,24,25]) enlargement of these fitter basins will outcompete lower fitness basins and create a regulation network that produces fit phenotypes more reliably. Although this result is somewhat underwhelming in this almost trivial (two attractor) system, in addition to the basic Hebbian principles, it also illustrates a further vital point. Specifically, the fact that the basin of attraction for the superior phenotypes is now almost 100% means that there are some initial conditions that used to lead natural selection of expression patterns to find the inferior phenotype but now evolution of expression patterns from these same initial conditions leads to the superior phenotype. That is, random variation in the expression pattern that would increase fitness by moving toward the inferior phenotype is being suppressed by the regulation network, and variation that moves the expression pattern toward the superior phenotype is being supported. This means that given the evolved regulation network, the evolutionary trajectory of the expression pattern is able to 'climb out' of the basin of attraction for the inferior phenotype and secure adaptation in the direction of the superior phenotype. Evolution of regulation that avoids sub-optimal phenotypes in a larger system is shown in Fig.7[1].

Ballistic development. Thus far the developmental network is only run for one time step (p=1) per application of natural selection. This is sufficient to induce significant correlations and redirect the evolutionary trajectory of expression patterns, as we have shown. But in general one might expect a regulation network to 'develop' an initial expression pattern into a fit adult expression pattern for many time steps without the need for selection to act on the result of every intermediate step. We therefore examine a 'ballistic' developmental trajectory (i.e. *run*(E,R) with p=5000, rather

[1] Here fitnesses are measured on thresholded expression values (>0→1, <0→-1) to ensure that an increase in fitness is the result of increasing the basin of attraction for a fit configuration pattern and not merely the result of increasing the magnitude of the expression levels (see measuring energy with the original weights rather than the learned weights [18]).

than 5000 iterations of the evolutionary cycle with p=1) using the regulation network evolved in Fig.7, applied to an initially random expression pattern. We find that even though selection is not being applied the fitness of the phenotype increases monotonically at each developmental step, and in fact the phenotypic attractor that is reached by this ballistic developmental process is the same attractor that is reached when selection was applied (Fig. 8). Thus selection on intermediate phenotypes (and epigenetic inheritance) has become redundant because development can now 'recall' the result of, or recapitulate, what was previously an entire evolutionary episode from any initial condition. Analogy with the Baldwin effect, where phenotypes that were previously acquired by lifetime learning are latterly exhibited innately [19], is provocative.

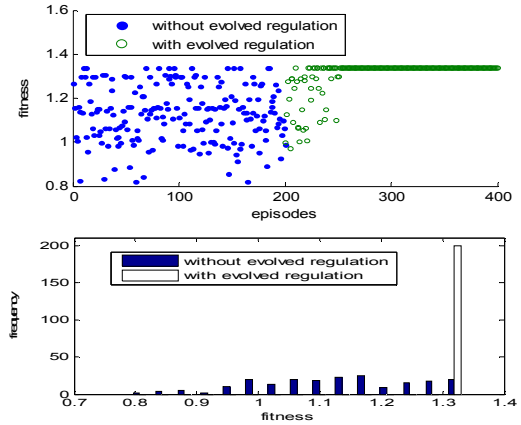

Fig. 7. As per Fig. 6 for a system of 30 genes with random epistasis in the fitness function (Eq.2 with each e_{ij} drawn randomly (-1,1)). The basin of attraction for the highest fitness optima is initially only 9.5%, meaning that 90.5% of episodes get stuck at some other sub-optimal phenotype. After the regulation network is evolved all of these inferior phenotypes are reliably evaded regardless of the initial gene expression pattern.

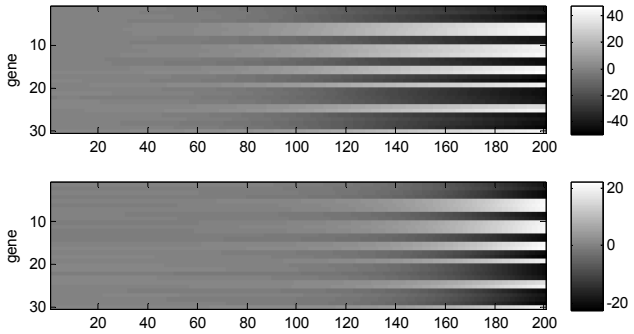

Fig. 8. 200 steps of an evolutionary episode with the evolved regulation network (upper) are accurately mimicked by ballistic (unselected) multi-step development using the same network (lower).

Discussion

Distal 'explanation'? On the one hand, the result of Fig. 7 is just what one might expect – selection favours fit phenotypes and if there are regulation networks that produce fit phenotypes reliably then they will be selected for. But this distal reasoning is misleading and obscures the proximal mechanism by which this result is produced. Note that a regulation network can preclude fit phenotypes just as easily, if not more so, than it might support them – it has 'masking' as well as 'guiding' possibilities [26] – and the evolution of a useful regulation network must not be taken for granted.

The point we illustrate in the initial results (Tables 1 & 2, Fig. 5) is that the evolved regulation network is not favouring fit phenotypes in a direct sense, it is merely canalising the *current* phenotype. This is not an obvious route to finding fit regulation networks and one might expect that, at best, it will ultimately result in canalising an average-fitness phenotype, *not* the fittest phenotype. But when the distribution of phenotypes visited over many evolutionary episodes has some correlations (or anti-correlations) that occur more frequently than others, it is these correlations that are ultimately reinforced by the regulation network (Fig. 6). If these correlations appropriately reflect the epistatic structure in the fitness landscape then they can enhance evolvability. In this manner the regulation network comes to represent the structure of the epistasis (or more exactly, the structure of the correlations between phenotypic characters produced by the epistasis) in the selective history over which the regulation network was evolved. But by the same reasoning, when the correlations in characters in the phenotypes visited do not reflect the epistatic structure of the fitness landscape in general, and instead reflect arbitrary phenotypic correlations, the regulation network will evolve to represent correlations that are not of especially high fitness. We demonstrate this by increasing the mutation rate on the regulation network, and/or increasing the duration of each evolutionary episode, such that the evolutionary history does not visit a representative sample of phenotypic attractors before the regulation network fixes on a particular attractor. On average this causes the regulation network to fix a phenotype with an average fitness rather than the highest fitness. Accordingly, it is not to be taken for granted that a gene regulation network will evolve to enhance high-fitness phenotypes just because such a network exists in the space of possible networks.

Proximal explanation. We should therefore investigate the proximal selection pressures involved in the initial result of Tables 2 & 3 (i.e. these data show that the selected changes to regulation connections are Hebbian but they do not explain why). Why is it that connections that reinforce the *current* phenotype are evolved instead of, say, connections that enlarge the basin of attraction for the fittest possible phenotype? (And how does this ultimately result in fit phenotypes?) To probe this issue we must consider the immediate selective gradients in the vicinity of the current phenotype. Specifically, for a change to a regulation connection to confer a selective advantage it must change the configuration of expression levels in a manner that increases fitness. However, most of the time, the current phenotype is a locally optimal configuration of gene expression levels. Thus, it might seem that the only way for a change to a connection to confer a fitness advantage would be when such a change moves the current phenotype out of the current local optimum and into a better one in a single mutation. But such a possibility is highly unlikely when the nearest phenotype of higher fitness is not an immediate neighbour.

In fact, something much more subtle is at work. Although most of the time the phenotype is *almost* locally optimal it is in fact constantly perturbed by the small environmental perturbations (line 2.b in Fig. 3). Changes to the regulation network can therefore be favoured by selection if they have the effect of returning the phenotype to the local optimum more quickly or more completely after this minor perturbation. In other words, we argue that changes to the regulation network are selected for merely because they make the current (almost locally optimal) phenotype more robust or more homeostatic. We test this hypothesis by removing line 2.a., the small environmental perturbations, and repeating the experiment shown in Table 2. In this case we find that there are no changes to the regulation network that are selected, in fact all changes are either neutral or deleterious. Thus the small environmental perturbations serve a dual role – they first provide (unregulated) phenotypic variation that selection can act on to find locally optimal phenotypes, but they also create instability in these phenotypes creating a selective gradient that favours a regulation network that canalises these phenotypes. We argue that this dual role of variation is not special to this particular model but will necessarily occur whenever random variation, necessary for evolution to act at all, is present.

From proximal causes to distal consequences. This proximal mechanism is also not very surprising given what one might expect from natural selection – if natural selection can act on the distribution of phenotypes in such a way as to narrow that distribution onto the fitter phenotypes, then a regulation network, for example, that provides such an outcome will be selected for. But canalisation – a *reduction* in the distribution of phenotypic characters – seems opposed to concepts of evolvability and *increases* in adaptability. However, a selection pressure for robustness can result in increased adaptability – in essence evolvability is the complement of canalisation [5]. The basic conceptual link is that restricting variation in phenotypic characters that are detrimental, whilst permitting continued variation in characters that have the potential to be beneficial, enhances adaptation rather than restricts it. But it is crucial to realise that in the current model the canalisation provided by the regulation network does not merely restrict variation in some characters but rather it reduces the degrees of freedom in the *correlation* of phenotypic characters [4].

In contrast, note that in Hinton and Nowlan's model [19] for example, canalisation acts to reduce the variation in each phene independently. This therefore cannot act like an associative memory – it is not a memory of what things have *co-occurred* (i.e. have been selected together in the same environments) only of what things have *occurred* (been selected). The fact that the memory in our evolved regulation networks is associative is evidenced by the fact that variation in all phenes is still possible (when the network canalises the fitter attractor it actually canalises both HHHH and LLLL). This is crucial because if no further variation in phenotypic characters was possible we would conclude that canalisation had precluded further adaptation, but when canalisation creates correlations in phenotypic variation it is plausible to interpret this as smarter adaptation, i.e. a more evolvable genotype, rather than an unevolvable genotype. This is really a matter of perspective however, since both types of canalisation (associative and non-associative) necessarily reduce the space of phenotypic possibilities.

Limitations and further work

Our gene expression network uses signed expression levels to facilitate straightforward comparison with Hebb's rule, but negative expression levels are biologically unnatural. We have also hinted at the sensitivity of the results to the timescales of evolutionary changes to expression patterns and to the regulation network, and to the period of the perturbations/ evolutionary episodes, but we have not yet examined this sensitivity carefully.

In related work we are interested in the question of whether individual agents in a complex adaptive system that can alter the strength of connections with one another will tend to do so in a Hebbian manner [17,27,28]. In this paper we have shown that selection on a network as a whole produces Hebbian changes to connections, but we suspect that the same effect occurs if each gene in the network is evolved independently. This hints at an explanation for how a network of 'selfish' genes can coordinate with one another in a manner that creates fit phenotypes despite being selected as individuals in sexual organisms. This then parallels work we are developing in the context of co-evolving species in an ecosystem where species may evolve the coefficients of a Lotka-Volterra system [27] or evolve symbiotic relationships [29], and connects with 'social niche construction' concepts [30].

The fact that natural selection is involved in this model should not to be mistaken for evidence of how 'clever' natural selection is. On the contrary, we have shown that given an appropriate (i.e. association-based) representation, a hill-climber can produce these results. Moreover, the proximal cause of these results is that selection is decreasing variability which is something that hardly warrants natural selection at all [17,18,31]. We think it more fruitful to ascribe the 'cleverness' of the result to the ability of an appropriate substrate to 'yield' or 'relax' to structured perturbation in a manner that reduces or dampens the effects of such perturbations [31]. This is supported by the observation that Hebbian changes to connections are equivalent to changes in connections that reduce the energy of a system [17].

Conclusions

Wagner *et al* [10] suggest that phenotypic correlations will evolve in a manner we recognise as Hebbian. Our conclusions, originating from separate motivations [11,17], agree but differ in emphasis – whereas Wagner *et al* address the rate of adaptation created by a correlated phenotypic distribution we emphasise the robustness or stability of a phenotype under environmental perturbation. But the mechanisms are deeply related because resilience is just another way to say that a phenotype 're-adapts' quickly. All of the other results we have shown – the enlargement of the basin of attraction for the current phenotype, the ability to 'recall' fit phenotypes that have been selected for in the past,

and the ability for a developmental trajectory to recapitulate what was previously an evolutionary trajectory – follow from this basic observation and dynamics that are already well-understood in neural networks. This theoretical framework helps us to better understand the relationship between homeostasis and evolvability (i.e. selection to differentially reduce variability facilitates structured variability), and shows that, in principle, a gene regulation network has the potential to exhibit 'recall' capabilities normally considered to be the exclusive purview of cognitive systems.

Acknowledgments. Chrisantha Fernando, Jason Noble, Alexandra Penn, Simon Powers, Marie Redron.

References

1. Bedau, M.A., McCaskill, J.S. *et al*., (2000) Open problems in artificial life, *Artificial Life*, 6(4):363-76)
2. Bedau, M.A. (2006) The Evolution of Complexity. Symposium on "The making up of organisms" Ecole Normale Supérieure, Paris.
3. Toussaint, M., & von Seelen, W. (2007) Complex adaptation and system structure, *BioSystems* 90: 769–782
4. Wagner G.P. & L. Altenberg, L. (1996). Complex adaptations and the evolution of evolvability. *Evolution* 50:967–976.
5. Kirchner, M. and Gerhart, J. (1998). Evolvability. Proc. Natl. Acad. Sci. USA, 95:8420–8427.
6. Izquierdo, E.J. & Fernando, C. (2008) The Evolution of Evolvability in Gene Transcription Networks. Proceedings of Alife XI.
7. Crombach A, Hogeweg P (2008) Evolution of Evolvability in Gene Regulatory Networks. *PLoS Comput Biol* 4(7).
8. Parter, M., Kashtan, N., Alon, U. (2008) Facilitated Variation: How Evolution Learns from Past Environments To Generalize to New Environments. *PLoS Comput Biol* 4(11): e1000206.
9. Lipson, H., J.B. Pollack, & N.P. Suh. (2002).. On the origin of modular variation. *Evolution* 56, 1549-1556
10. Wagner, G.P., Pavlicev, M., & Cheverud, J.M. (2007) The road to modularity. *Nature Reviews Genetics* 8: 921-931.
11. Mills, R. (2010) *How Micro-Evolution Can Guide Macro-Evolution: Multi-Scale Search via Evolved Modular Variation.* PhD thesis, ECS, University of Southampton.
12. Pavlicev, M, Kenney-Hunt, J.P., Norgard, E.A., Roseman, C.C., Wolf, J.B. & Cheverud, J.M. (2008) Genetic variation in pleiotropy: differential epistasis as a source of variation in the allometric relationship between long bone lengths and body weight. *Evolution.* 2008 Jan;62(1):199-213.
13. Hebb, D.O. (1949) *The organization of behaviour*. New York: Wiley.
14. Ackley, D.H., Hinton, G.E., Sejnowski, T.J. (1985) A Learning Algorithm for Boltzmann Machines, *Cognitive Science*, 9: 147-169.
15. Hopfield, J.J. (1982) Neural networks and physical systems with emergent collective computational abilities, *PNAS USA*, 79 (8) 2554.
16. Fernando, C. Liekens, A.M.L., Bingle, L.E.H., Beck, C., Lenser, T., Stekel, D.J. ,Rowe, J.E. (2008) Molecular circuits for associative learning in single-celled organisms. *Journal of the Royal Society Interface.* 6:463-469.
17. Watson, R. A., Mills, R, and Buckley, C. L. (2009) *Global Adaptation in Networks of Selfish Components: Emergent Associative Memory at the System Scale.* Submitted/Tech. Report, ECS, University of Southampton.
18. Watson R.A., Buckley, C.L. & Mills, R. (2010). Optimisation in 'Self-modelling' Complex Adaptive Systems, *Complexity* (under revision)/ Tech. Report, ECS, University of Southampton.
19. Hinton, G.E. & Nowlan, S.J. (1987) How Learning Can Guide Evolution. *Complex Systems*, 1, 495-502
20. Ashby, W.R. (1952) *Design for a Brain*, Chapman & Hall
21. Rechenberg, I., *Evolutionsstrategie: Optimierung technischer Systeme und Prinzipien der biologischen Evolution* , Frommann-Holzboog, Stuttgart, 1973.
22. Beer, R.D. (1995) On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*. 3(4).
23. Kauffman, S.A. (1989) Adaptation on rugged fitness landscapes. *Lectures in the Sciences of Complexity*, 527-618.
24. Fontanari, J.F. (1990) Generalization in a Hopfield network. *Journal de Physique*, 51, 2421-2430.
25. Kryzhanovsky, B. & Kryzhanovsky, V. (2008) Binary Optimization: On the Probability of a Local Minimum Detection in Random Search. *ICAISC 2008*: 89-100.
26. Mayley, G. (1997). Guiding or hiding: Explorations into the effects of learning on the rate of evolution. *Procs. ECAL 4*, 135-144.
27. Lewis, M. (2009) *An Investigation Into The Evolution Of Relationships Between Species In An Ecosystem*. MSc Dissertation, ECS, University of Southampton.
28. Davies, A.P., Watson, R.A., Mills, R., Buckley, C. L., Noble, J. (2010) 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. *ALife XII* (to appear).
29. Watson, R.A., Palmius, N., Mills, R., Powers, S.T., & Penn, A.S. (2009a) Can Selfish Symbioses Effect Higher-level Selection? *European Conference on Artificial Life 2009*. (in press).
30. Powers, S.T., Mills, R. Penn, A.S., Watson, R.A. (2009) Social niche construction provides an adaptive explanation for new levels of individuality (ABSTRACT), *Procs of Workshop on Levels of Selection and Individuality in Evolution*, *ECAL 2009*.
31. Watson, R.A., Mills, R. Buckley, C.L., Powers, S.T., Penn, A.S., Davies, A.P., Noble, J. & Bullock, S.G. (2010) Adaptation without natural selection (ABSTRACT). *ALife XII* (to appear).