The evolution of phenotypic correlations and "developmental memory"

Richard A. Watson[1], Günter P. Wagner[2], Mihaela Pavlicev[3], Daniel M. Weinreich[4], Rob Mills[5]

[1]Natural Systems Group, ECS/Institute for Life Sciences/Institute for Complex Systems Simulation, University of Southampton. SO17 1BJ. U.K.. R.A.Watson@soton.ac.uk

[2]Ecology and Evolutionary Biology, Yale University. CT 06477. USA. gunter.wagner@yale.edu

[3]Theoretical Biology, University of Vienna, Austria. mihaela.pavlicev@univie.ac.at

[4]Ecology and Evolutionary Biology/Center for Computational Molecular Biology, Brown University, RI 02912. USA. daniel_weinreich@brown.edu

[5]ECS, University of Southampton. SO17 1BJ. U.K. Rob.Mills@soton.ac.uk

Development introduces structured correlations among traits that may constrain or bias the distribution of phenotypes produced. Moreover, when suitable heritable variation exists, natural selection may alter such constraints and correlations, affecting the phenotypic variation available to subsequent selection. However, exactly how the distribution of phenotypes produced by complex developmental systems can be shaped by past selective environments is poorly understood. Here we investigate the evolution of a network of recurrent non-linear ontogenetic interactions, such as a gene regulation network, in various selective scenarios. We find that evolved networks of this type can exhibit several phenomena that are familiar in cognitive learning systems. These include formation of a distributed associative memory that can 'store' and 'recall' multiple phenotypes that have been selected in the past, recreate complete adult phenotypic patterns accurately from partial or corrupted embryonic phenotypes, and 'generalise' (by exploiting evolved developmental modules) to produce new combinations of phenotypic features. We show that these surprising behaviours follow from an equivalence between the action of natural selection on phenotypic correlations and associative learning, well-understood in the context of neural networks. This helps to explain how development facilitates the evolution of high-fitness phenotypes and how this ability changes over evolutionary time.

Keywords: associative learning, evolvability, evo-devo, adaptation

Introduction

The extraordinary ability of natural selection to adapt organisms to diverse and challenging environments depends fundamentally on the supply of appropriate heritable phenotypic variation. The distribution of phenotypic variants that occur as a result of genetic and environmental variation is shaped by developmental processes that transform the embryonic phenotype into the adult form. These developmental processes involve complex interactions that can introduce correlations between phenotypic traits, causing some traits to co-vary, creating patterns of phenotypic variation that are thereby partially non-random. Since developmental processes are themselves a product of evolution, such biases and constraints can, in principle, be shaped by past selection (Riedl 1978, Raff 2000, Wagner, Pavlicev & Cheverud 2007, Izquierdo & Fernando 2008, Hendrikse, Parsons & Hallgrímsson 2007, Pavlicev & Wagner 2012, Crombach and Hogeweg 2008, Clune et al. 2013, Brigandt 2007, Draghi et al. 2010).

We seek general organisational principles to understand how past selective environments can alter phenotypic correlations and hence shape the distribution of phenotypic variants produced by development in adaptive ways (Toussaint & von Seelen 2007, Wagner, Pavlicev & Cheverud 2007). In particular, we are interested in the idea that developmental processes, shaped by past selection, may constitute a 'memory' of phenotypes or phenotypic features that have been selected for in the past. Such a *developmental memory* would cause development to be predisposed to produce these phenotypic features in subsequent evolution. To the extent that future selective environments have properties that are similar to past selective environments, such a developmental memory could enrich variation for well-adapted phenotypes.

Mechanistically, phenotypic correlations in natural organisms arise in a number of different ways from interference between expression pathways, or transcription factors in a gene regulation network (GRN) (effecting correlated or anti-correlated gene activity levels), to the physiological interactions involved in macro-scale morphological growth. Heritable genetic variation affecting phenotypic correlations has been shown in quantitative data from many organisms (Cheverud et al. 2004, Pavlicev et al. 2008, Leamy, Pomp & Lightfoot 2009, Chevillon et al 1997, Lenski 1988a/b, Kim Huh & Fay 2009). This means that phenotypic correlations can change as a result of evolution by natural selection (Delph et al 2011). Examples have been documented with respect to fore and hindlimb correlations in mammals (Young et al., 2005) and in primates in particular (Young, et al.. 2010). Characterising how these interactions change over evolutionary time is crucial to understanding the properties of developmental processes and how particular phenotypic patterns can be preferentially expressed (Guillaume & Otto 2012). To begin to explain these patterns Pavlicev, Cheverud & Wagner (2011) provide a detailed analysis of the direct selective pressure on relationship loci (rQTL) affecting associations between two quantitative traits (Pavlicev et al. 2008). They show that selection modifies the sign and magnitude of the correlation in the direction that increases phenotypic variation in the direction of selection, hence increasing the rate of evolutionary change. Notice that, on the assumption that developmental constraints evolve slowly compared to the quantitative traits they affect, these developmental correlations will bias the combinations of trait values that will be produced in future, and in particular, will bias them to reproduce trait combinations that have been selected for in the past.

In larger developmental systems, the application of these simple selective pressures could support the emergence of developmental modularity, i.e., cause subsets of traits to vary together but independently of other subsets (Wagner, Pavlicev & Cheverud 2007). Early work on this idea (Lipson, Pollack & Suh 2002) represented the genotype-to-phenotype (G-P) mapping with a matrix that implemented a linear transformation from environmental 'input' to phenotypic 'output', directly representing phenotypic correlations, and demonstrated that this matrix can evolve to allow elements of the phenotype to vary independently if the variation in the selective environment favours such independence. Kashtan et al. (2009) also use a matrix-based representation of linear correlations and find that modularity in the evolved matrix corresponds to the correlations within the input/output vectors.

These works utilise a linear statistical model of phenotypic correlations and this has some limitations in the type of distributions it can represent. In particular, a linear model can represent a phenotypic distribution that is directionally biased (e.g., creating a genetic 'line of least resistance' for the evolution of a population, Schluter 1996) but it cannot, for example, model a phenotypic distribution that is multi-modal. For example, if a developmental process can produce both sepals and petals, a linear model can capture the correlations amongst the multiple features of sepals and petals, but the resultant distribution also includes phenotypes all along the (multi-dimensional) line in between sepals and petals.

In principle, a non-linear model of a developmental process or a more complex G-P mapping could represent and/or produce multi-modal phenotype distributions. For example, Kashtan, Noor & Alon

(2007, also Kashtan & Alon 2005) evolve phenotypes (logic functions and RNA secondary structures) using a genotype space that allows small changes in genotype to 'switch' between previously selected phenotypes that are far apart in feature space. Relatedly, Parter, Kashtan & Alon (2008) develop analogies with memory, i.e., past selection for a phenotype changes the G-P mapping such that development more readily produces that phenotype in subsequent evolution. Notably, they also illustrate an ability for evolved genotypes to "generalise to future environments, exhibiting high adaptability to novel goals", showing that the memory is not merely reproducing past phenotypes in a one-to-one (or 'rote learning') manner. But is memory just a loose analogy for the fact that the G-P map has been altered by past selection? And how can we use what we know about the selective pressures acting on correlations in the simple linear models to understand the capabilities and limitations of more sophisticated non-linear developmental processes?

Here we build on this prior work with the aim of identifying organisational principles to predict how past selection shapes the properties of non-linear developmental processes. Rather than assuming a simple linear model of phenotypic correlations, or a highly complex G-P mapping where we would have limited insight into intrinsic biases, we assume a developmental model that is capable of exhibiting sophisticated behaviours yet simultaneously simple enough to understand exactly how it becomes altered as a function of past selective environments. Our G-P mapping is defined by an interaction matrix, as common in prior work, but here this matrix represents a set of ontogenetic interaction coefficients for a simple but recurrent and non-linear model of ontogenetic interactions. Some non-linearity in the mapping is important (as we will show) but this can be of a simple, natural form; e.g., a simple sigmoid function characteristic of many natural systems where the effect of a forcing variable attenuates in the extremes (e.g., the effect of a transcription factor saturates at high concentrations, see *Methods*). These recurrent non-linear interactions transform a set of embryonic phenotypic characters into their adult form over multiple developmental time steps.

We find that this non-linear G-P mapping can exhibit a developmental memory capable of preferentially producing multiple distinct phenotypes that have been selected for in the past (i.e., a multi-modal phenotypic distribution). Such a developmental process can also recreate a complete adult phenotypic pattern accurately from an embryonic phenotype that partially resembles a previously selected phenotype. It can also show the capability to generalise past phenotypic patterns, e.g., by evolving developmental modules and producing new phenotypes that are built from novel combinations of those modules.

In addition to illustrating these memory behaviours in a simple model of development, the main contribution of this work is that we also show that there is an existing theoretical framework that can be transferred from another discipline to understand these surprising capabilities. This builds on two previous observations. First, the dynamical and functional capabilities of gene networks and neural networks are mathematically equivalent (Vohradsky 2001a,b). Both have state dynamics controlled by a non-linear weighted sum of interactions between state variables. Thus for any given neural network there exists a gene network capable of computing the same functions or exhibiting the same dynamical behaviours. However, this observation does not address how gene networks change over time by evolution nor how neural networks change over time by learning. Evolved changes to ontogenetic interactions are the result of random variation and selection whereas neural connections can be changed by purpose-specific learning mechanisms that alter synaptic connections in a directed fashion. Thus although, in principle, there exist gene networks that can exhibit the behaviours of any neural network, there is no obvious reason to believe that evolution will be able to find GRNs that behave like well-trained neural networks that produce interesting behaviours.

Second, Pavlicev, Cheverud & Wagner (2011) show that if two traits are both under positive directional selection (i.e., for increasing trait values) or both are under negative directional selection, then selection

favours an increase in correlation (or decrease in anti-correlation) of those traits. Conversely, if the traits are selected contrariwise, i.e. one is under positive directional selection and the other negative, then selection favours anti-correlation (or decreases in positive correlation) of those traits.

A new insight links these previous observations in a surprising and productive way. Specifically, we note that the selective pressures on phenotypic correlations observed by Pavlicev, Cheverud & Wagner are equivalent to a simple and well-known neural learning rule. *Hebbian learning* (Hebb 1949), is a simple associative learning mechanism well understood in the context of connectionist models of memory and knowledge representation (Ackley, Hinton & Sejnowski 1985, Rumelhart et al. 1986, O'Reilly & Munakata 2000)[1]. *Hebb's rule* simply states that the change in strength of a synaptic connection, Δw_{ij}, is proportional to the co-activation of the neurons it connects: $\Delta w_{ij}=rs_is_j$, where $r>0$ is a learning rate, and s_x is the activation level of node x in response to a training pattern. This type of learning is often paraphrased as '*neurons that fire together wire together*'. We note that Pavlicev, Cheverud & Wagner's observation tells us that, in effect, 'traits that are selected together correlate together'. That is, the direction of selective pressures on individual relational loci described above (also matching observations in Kashtan et al. 2009, Watson et al 2010a, and in the present models) has the same relationship with a selective environment that the direction of changes to synaptic connections in a learning neural network has with a training pattern. In other words, gene networks *evolve* like neural networks *learn*.

Bringing together these two observations with this new insight explains the memory behaviours we observe in an evolved network of recurrent non-linear interactions. That is, a gene network can evolve regulatory interactions that 'internalise' a model of past selective environments in just the same way that a learning neural network can store, recall, recognise and generalise a set of training patterns. Recognising this equivalence between associative learning and evolution of phenotypic correlations thus provides access to an established theoretical framework that we can use to characterise organisational principles describing how regulatory interactions evolve; including the affordances and limitations of developmental memory, the minimal conditions for such behaviours and their potential impact on evolvability. It is a main aim of this paper to explain the dynamical and functional equivalence of these evolved developmental behaviours to the capabilities that are already well-defined and understood in learning cognitive models, and in the experiments that follow we will explain and illustrate each of these behaviours.

The Model

Representation of individuals and developmental genotype-phenotype mapping

The phenotype of an individual at developmental time step, t, is described by a set of N phenotypic characters or traits, naturally represented by a vector, $P(t)=\langle p_1(t), p_2(t),\ldots p_N(t)\rangle$, $p \in \mathbb{R}$. The genotype of an individual has two parts, naturally represented by a vector of direct effects on traits, $G=\langle g_1, g_2,\ldots g_N\rangle$, and the elements b_{ij} of an interaction matrix, B (Wagner 1989; Lipson, Pollack & Suh 2002, Jones, Arnold & Bürger 2007, Lande & Arnold 1983, Kashtan et al. 2009). Whereas previous work utilised a linear model of phenotypic correlations where B^T was a genetic covariance matrix among the adult characters, here we analyse a non-linear model of a developmental process where b_{ij} represents the

[1] Other work has investigated the potential to implement associative learning mechanisms in various non-neural systems (e.g., metabolic networks - Fernando et al. 2008); and investigated the ability of evolution to find such mechanisms (McGregor et al. 2012) that can then operate within the lifetime of the individual. Here we do not select for an associative learning mechanism, we simply evolve developmental interactions.

interaction coefficient between trait *i* and trait *j* within a dynamical ontogenetic process modelled as follows.

Let the initial embryonic phenotype of an organism at developmental time *t*=0, i.e. *P*(0), be *P*(0)=*G*, where the components of *G* are the direct effects on the embryonic phenotype. In each developmental time step thereafter, the phenotype vector is updated by a non-linear transform determined by the matrix *B*, i.e. the weighted influence of each character on each other character, and a decay term. Specifically, in linear algebra form,

$$P(t+1) = P(t) + \tau_1 \sigma\big(B \times P(t)\big) - \tau_2 P(t), \qquad \text{Eq. 1.}$$

where τ_1=1 is a rate constant controlling the magnitude of the interaction terms, τ_2=0.2 is a decay rate, and σ is a sigmoidal function (applied to each element of the vector) that non-linearly limits the influence of interactions, we choose, $\sigma(x)=\tanh(x)$. This contrasts with a linear mapping where $\sigma(x)=x$ (for a single-step G-P mapping where $P(0)=G$, this simplifies to $P = G + \tau_1(B \times G) - \tau_2 G$, or with suitable adjustment in the elements of *B*, simply $P = B \times G$, as per, for example, Lipson, Pollack & Suh (2002) and Kashtan et al. (2009)). In the absence of interactions, that is if all off-diagonal elements of *B* are zero, Eq. 1 assumes a decreasing rate of change in the phenotype states as the size of the characters increase. The reason is that the sigma function limits the size of the growth increments and thus the relative change is decreasing over time. The off-diagonal elements of *B* introduce interdependencies among the different characters, where the size of one trait influences how much another trait grows at any time step.

This class of non-linear transformation has analogues in many natural systems and biological processes. An example of ontogenetic process that can be modelled this way is a gene-regulation network where *P* is a pattern of gene activity levels (deviations from mean levels), and *B* is a network of up- and down-regulatory interactions (Wessels, van Someren & Reinders 2001), which develops the 'embryonic' activity levels of each gene into an 'adult' pattern of activity. As per Eq. 1, the new activity level of a gene, $p_i \in P$, in a single time step is given by:

$$p_i(t+1) = p_i(t) + \tau_1 \sigma\left(\sum_{j=0}^{n} b_{ij} p_j(t)\right) - \tau_2 p_i(t), \qquad \text{Eq. 2}$$

where b_{ij} is the regulatory effect of gene *j* on the activity level of gene *i*, and $\sigma(x)=\tanh(x)$ as before, represents a non-linear interaction effect (e.g., saturation of a transcription factor). This regulatory network (Eq.2) is simply a specific interpretation of our general ontogenetic process (Eq.1), written in the form more common to the GRN literature (this is the general form for continuous outputs rather than the simplified form for binary-outputs, where σ would be a threshold function). In this example, we refer to each gene expression level, p_i, as a phenotypic character or trait, and the gene expression profile, *P*, as the phenotype.

The fitness of the organism is determined by the adult phenotype, *P**, after a fixed number of developmental time steps, *T*, i.e. *P**=*P*(*t*=*T*). This vector of gene expression levels may be interpreted as an attractor of a regulatory network corresponding to a cell type (Kauffman 1993, Huang et al. 2005), for example, or the product of a developmental process at more physiological scales. The adult phenotype is thus the result of the development of an embryonic phenotype (given by *G*) governed by the interaction matrix *B* over *T* time steps, i.e., *P**=*develop*(*G*,*B*,*T*).

Note that, in general, interaction terms can have 'mean' and 'relational' effects – that is, they can have an effect on the mean value of a phenotypic character that is produced by development as well as on the correlation between phenotypic characters (see Supplementary text (a)).

Evolutionary model

We model the evolution of segregating alleles in *G* and *B*, and consequent changes in the mean phenotypic traits of a population, over multiple generations using numerical simulation.

In this model, the fact that sexual recombination reduces or removes linkage between interaction alleles and direct alleles does not prevent natural selection on evolvability (Sniegowski & Murphy 2006) because the mean effects of mutations in *B* will cause those alleles to change in frequency even when there is no linkage with loci in *G*. To emphasise this point here we show that strong selection weak mutation (SSWM) assumptions (Gillespie 1984), i.e. when each new mutation is fixed or lost before the next new mutation occurs, such that only one locus at a time is polymorphic, are sufficient to produce the effects shown. The evolution of sexual and asexual populations are equivalent under these assumptions. Under SSWM the population is straightforwardly represented by a single genotype ($\bar{G}$ and $\bar{B}$) representing the population mean genotype. This genotype uniquely determines an adult phenotype, $\overline{P^*}$, via the developmental model, representing the population mean phenotype. For clarity of exposition we model haploid genotypes undergoing point mutations (see below). These mutations include mutations to *G*, the direct effects on embryonic quantitative traits, and mutations altering the interaction matrix *B*. Under these conditions, a simple hill-climbing model of selection is sufficient (Supplementary text (b)) and avoids obfuscating what is essentially a very simple and obvious selective pressure acting directly on the interaction coefficients. Each 'generation' of our simulation thus corresponds to multiple generations of a natural population during which the fixation or loss of a new mutation occurs. For investigations incorporating polymorphic populations and sexual recombination, and for analysis of selection pressures due to neutral relational alleles, see the linear models of Pavlicev, Cheverud & Wagner (2011).

Selection and varying selective environments

The fitness of an adult phenotype (in a single selective environment) is determined by constant directional selection on each phenotypic character defined by a selective environment, $S=\langle s_1,s_2,s_3,\ldots s_N\rangle$, $s \in \{-1,1\}$, such that fitness increases with alignment of the phenotype with this 'target' phenotype vector. Specifically, the fitness, *w*, of a phenotype is calculated from the scalar product of P^* and *S*, i.e. $w(P^*) = 1 + P^* \cdot S$. The elements of *S* thus determine the direction of selection on each of the phenotypic characters (Pavlicev, Cheverud & Wagner 2011).

Lipson, Pollack & Suh (2002), Kashtan & Alon (2005), Kashtan et al. (2009), Kashtan, Noor & Alon (2007), Draghi & Wagner (2008), Draghi & Wagner (2009) and Pavlicev, Cheverud & Wagner (2011) all investigate evolution in varying selective environments. This makes intuitive sense because stabilising selection in a single environment offers no advantage to genotypes that admit variability, structured or otherwise, and will simply tend to canalise phenotypes (Waddington & Robertson 1966) reducing phenotypic variability in all dimensions. In contrast, directional selection can favour genotypes that *increase* variability in the direction of selection. Moreover, rather than simply decreasing variability in some traits and/or increasing it in others, selection over a distribution of selective environments admits the possibility that a well-adapted developmental process may reflect correlations or deeper structural properties of this distribution.

In the experiments utilising variable environments, a different target phenotype is selected uniformly at random from a small set of phenotypic targets (described below for each experiment) after every 2000 evolutionary time steps. Environmental change is assumed to be slow compared to the generation time of

the population such that a given individual derives its fitness from only one selective environment. This means that phenotypic plasticity is not selected for explicitly since an individual only needs to exhibit one phenotype to maximise fitness.

Assessing the phenotypic correlations caused by development

In our investigations we observe the evolved values of G and B, and the adult phenotypes, P^*, produced. We also observe the pre-disposition of a given developmental network, defined by B, to produce particular adult phenotypes. To observe the distribution of adult phenotypes produced by a given B, we mutate G (or, in some cases, artificially manipulate G), and observe the phenotypes produced without applying further selection. In the limit where the genetic information in G is completely destroyed by mutation, such that $P(0)$ is an unbiased distribution of phenotypic patterns, each resultant phenotype indicates the intrinsic propensity of development to produce a particular combination of adult phenotypic characters. Observing the distribution of adult phenotypes produced by this method (without further selection) provides a direct way of assessing the developmental memory contained in B.

Model parameters

The number of developmental time steps, T, is set at T =10. All values in G and B are initialised to zero. Mutation on G is applied every evolutionary time step by adding μ_1, drawn uniformly in the range +/- 0.1, to a single trait (selected uniformly at random). The magnitude of direct effects is capped at +/-1, i.e., $|g| \leq 1$. A meaningful sense of developmental memory requires that the evolved characteristics of developmental constraints must be more slow-changing than the phenotypic variation they control. We therefore assume that the amount of heritable variation affecting correlations, B, is significantly lower than that affecting direct effects, G. Our simulations are sensitive to the relative ratio of mutation on G and B, but only in the sense that the latter needs to be sufficiently small (see *Analysis).* In our simulations with N=8 we find it sufficient that the rate and magnitude of mutations applied to B are 1/15th of those applied to G. Thus, in each evolutionary time step, mutation on one entry in B (selected uniformly at random) is applied with probability 0.067, by adding μ_2 in range +/- 0.0067. Simulation of large N is handled differently (see Experiments 3 and 4).

Experiments and Results

Experiments are conducted in a number of different evolutionary scenarios designed to directly investigate the tendency of evolved developmental processes to exhibit specific memory capabilities.

Experiment 1) Single selective environment

Experiment 1 assesses the basic effect of selection on interaction coefficients as a function of a single selective environment. In Fig. 1 (a-b) we observe that interaction coefficients divide into two classes; when traits are selected together, either ++ or --, this produces a selective pressure for positive b_{ij}, and when traits are selected contrariwise, either -+ or +-, this selects for negative b_{ij}. This agrees with previous observations in linear models (Pavlicev, Cheverud & Wagner 2011), i.e. direction of selection on b_{ij} is the same as the sign of $s_i s_j$. The reason for this is not difficult to understand intuitively. Natural selection favours any change that moves a phenotypic character in the direction of selection. If heritable variation in regulatory interactions produces a mean effect on phenotype this will be utilised, and the sign of that change will depend on the sign of the direction of selection on a given trait, and also the sign of the other character – either pushing it in the same direction or the opposite. Thus whereas selection on direct effects reacts to the direction of selection on individual traits, selection on phenotypic correlations naturally responds to the correlation between the directions of selection on two traits. Here we emphasise

that this is also in agreement with Hebb's rule, $\Delta w_{ij}=rs_is_j$ ($r>0$) (see *Analysis*). Fig. 1.b shows the evolved interactions whereas Fig. 1.c shows the interaction coefficients derived directly from Hebb's rule. This shows complete agreement in the sign of evolved interactions with Hebbian principles.

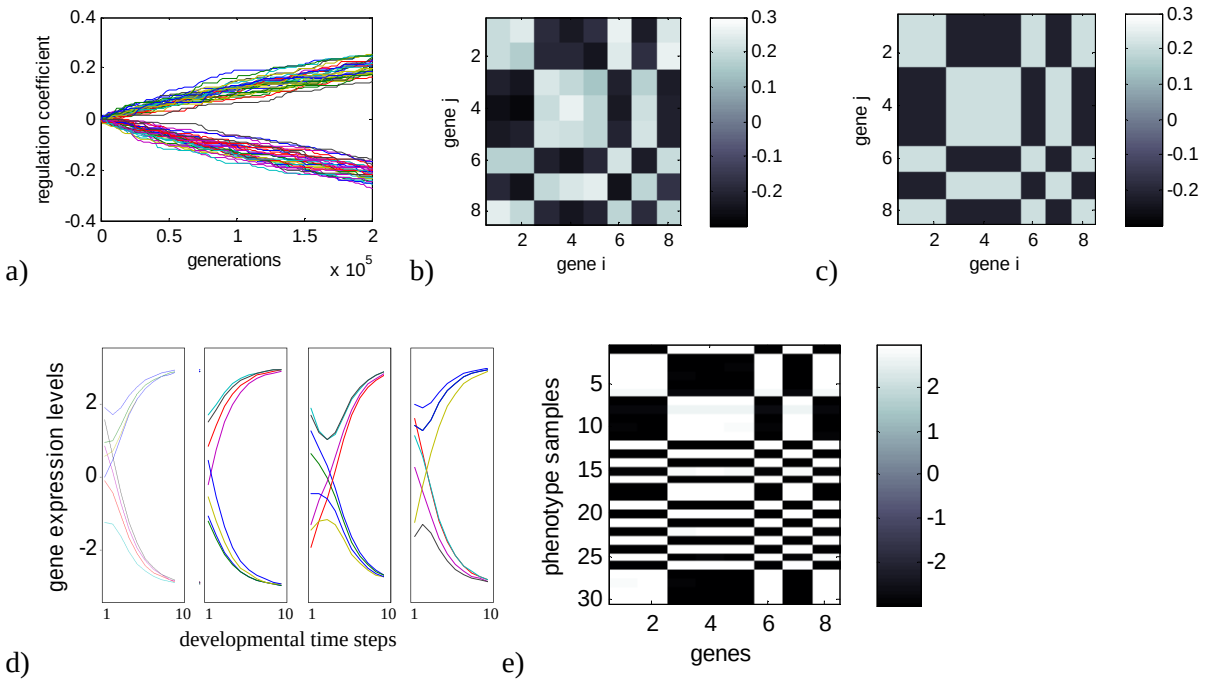

Fig. 1: Experiment 1. A system of 8 genes, selecting for a single (arbitrary) phenotypic pattern (*S1*=++---+-+). **a**) Regulatory interaction coefficients evolve into positive and negative classes. **b**) The matrix of evolved regulatory interactions in the gene regulation network (at generation $2{\cdot}10^5$). We observe that the selection pressure on interaction coefficients is such that self-interactions (b_{ij}, $i=j$) increase in all cases, other interaction coefficients increase approximately symmetrically ($b_{ij} \approx b_{ji}$), and more specifically, b_{ij} has the same sign as s_is_j. **c**) Interaction matrix derived from Hebb's rule (i.e., $w_{ij}=rs_is_j$) rather than evolution (r is scaled to match the same average magnitude as B). Note that the pattern of positive and negative values correctly predicts the pattern in the evolved regulatory interactions. **d**) Gene expression levels over developmental time for four independent developmental trajectories (from random G). All expression levels are approximately saturated, either at positive or negative extremes (deviations from normal), in the adult expression pattern. **e**) 30 independent adult phenotypes developed from random G. Development produces either the target phenotype or its complement.

Supplementary Text (c) confirms that the effect of these evolved developmental interactions on the distribution of adult phenotypes is to align phenotypic variation with the direction of selection as predicted (Pavlicev, Cheverud & Wagner 2011). Fig. 1 (d-e) shows that the sign of phenotypic characters can change during development to reproduce phenotypic patterns that correspond to the previously selected phenotype. If genetic information in G is sufficiently depleted, evolved interactions also produce the complement of the target pattern because B only controls the correlations and not the signs of each trait. Although this experiment demonstrates a bias to produce a previously selected phenotype, it is only remembering *one* pattern – in this sense it is an impoverished demonstration of 'memory', or more exactly, it is demonstrating little more than simple canalisation.

Experiments 2 & 3) Varying selective environment: Multiple memories

Experiment 2 evolves a regulatory network in a varying selective environment using two target patterns. This tests the propensity of a developmental process to accurately reproduce an adult phenotype including features that are contradicted by another target pattern. Fig. 2.a shows that evolved interactions fall into three classes. Some evolve at a constant positive rate; these arise from pairs of traits that are positively correlated in *both* patterns (e.g., genes 2 and 6 are ++ in *S1* and -- in *S2*), likewise negative interactions

evolve at a constant rate between pairs of traits that have opposite signs in *both* patterns. When the correlation of a pair of traits in one pattern is contradicted by the correlation of that pair in the other pattern (e.g., $s_1 s_2 > 0$ in *S1* and $s_1 s_2 < 0$ in *S2*) the corresponding regulatory interactions (e.g., b_{12} and b_{21}) are unable to record the correlation of either target pattern and remain near zero on average (Kashtan et al (2009) observe the same phenomenon). Fig. 2.b-c shows that, again, the classes of values (positive, negative and near zero) in the evolved interaction coefficients agree exactly with the pattern of interaction coefficients derived directly from Hebb's rule.

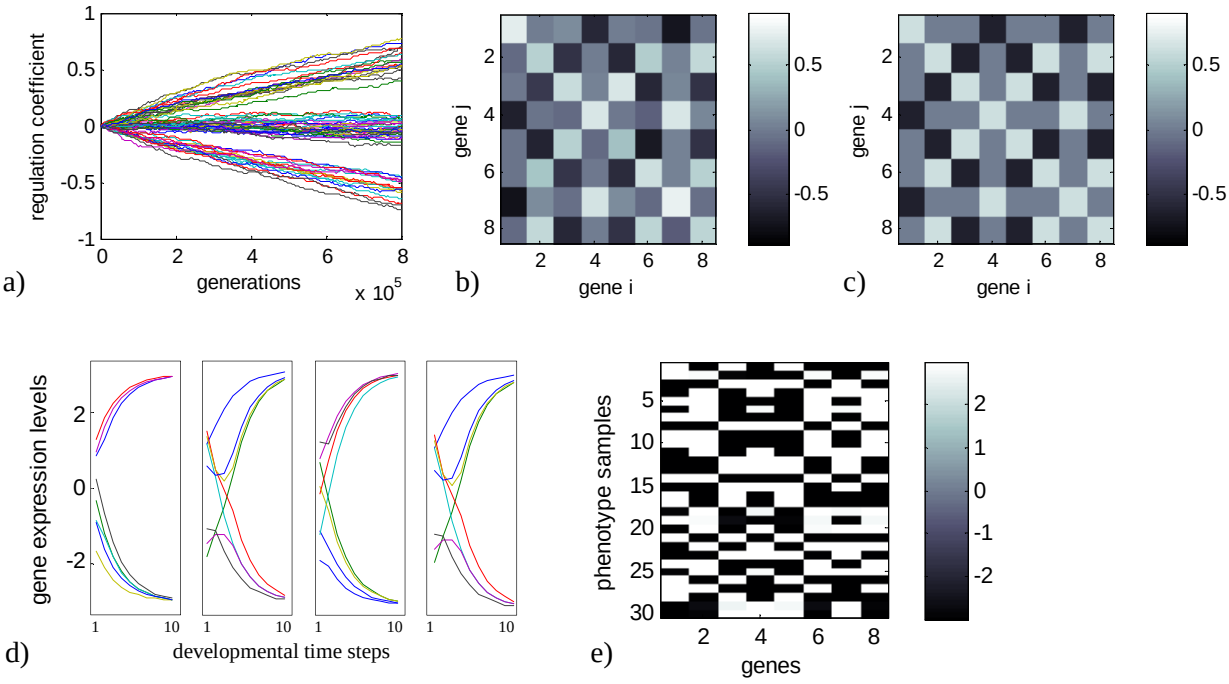

Fig. 2: Experiment 2. Evolved interaction coefficients for changing environment (two target patterns *S1*=++---+-+, *S2*=+-+-+---). Target patterns alternate each 2000 generations. **a**) Interaction coefficients over evolutionary time fall into three classes. **b**) Evolved interaction matrix after $8 \cdot 10^5$ generations. **c**) Interaction matrix derived from Hebb's rule (summed over the two patterns) showing that the pattern of positive/negative/zero values in the evolved interaction coefficients matches exactly the pattern predicted by Hebb's rule. **d**) Gene expression levels over developmental time for four independent developmental trajectories (from random *G*). **e**) 30 example adult phenotypes (one per row) developed using evolved interaction network. Development produces either one of the target phenotypes or their complements.

The particular two target phenotypes used in Experiment 2 (Fig. 2) are *S1*=++---+-+, *S2*=+-+-+---; These are arbitrary bit patterns but in an organism they may represent the particular pattern of gene expression corresponding to two different phenotypes. In this small example (N=8) we can see easily that the evolved interactions agree with the interactions derived from Hebbian learning (Figs. 2.b-c), and that the effect of these evolved correlations is that development reproduces either of the two target patterns (Fig. 2.e). Experiment 3 (Fig. 3) tests the generality and robustness of this result (and several additional effects) in a much larger genetic system. Using phenotypes that represent a recognisable image (rather than an arbitrary bit pattern) makes it intuitively easy to interpret whether a particular phenotype is being produced or not. We use an example where *S1* is an image of a particular well-known phenotype and *S2* is an image of Donald Hebb, the neuropsychologist from whom Hebbian learning takes its name (Fig. 3. a-b).

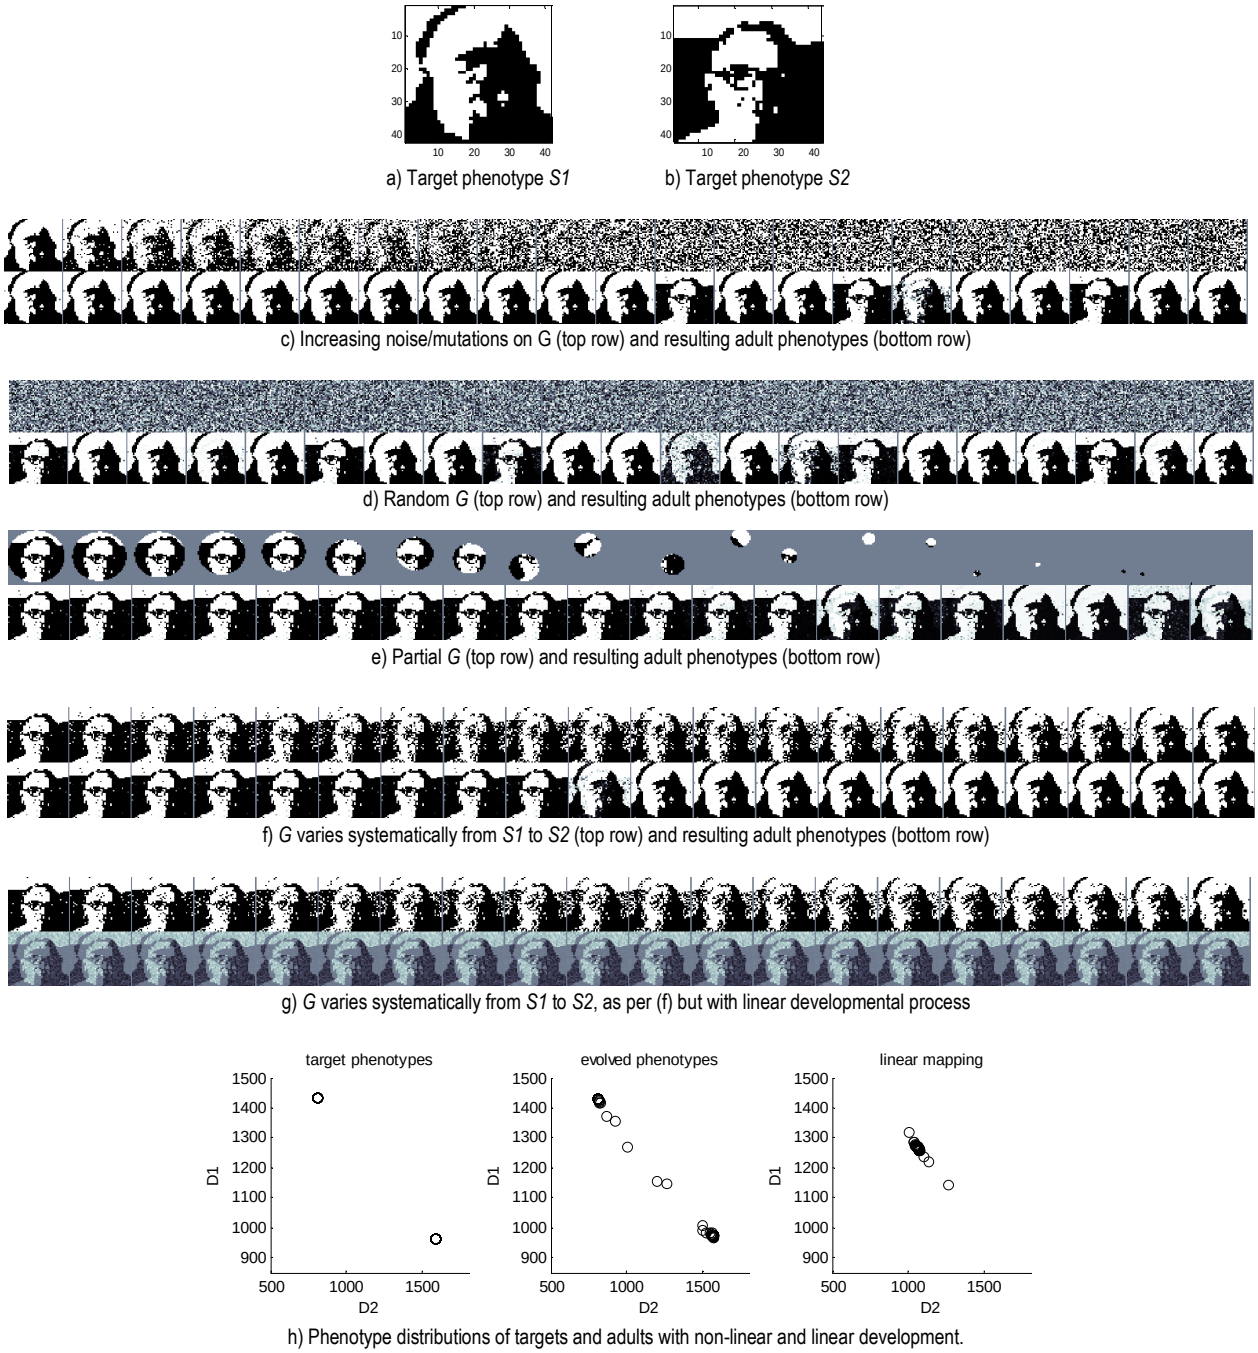

Fig 3. Experiment 3. a) Target phenotype *S1*=image of Charles Darwin; each pixel indicates whether a given gene should be up-regulated (white) or down-regulated (black) in the target phenotype. b) Target phenotype *S2*=image of Donald Hebb. c) Mutations applied to the direct-effects of the genotype, *G*, increasing levels of random mutations from 0 to 100% (replacing 5% of the elements of *G* with random alleles, +/-1, in each step left to right) starting from target *S1*. The corresponding adult phenotypes in each case, developed using evolved interaction network, *B*, show high robustness. d) Phenotypes produced from random *G* illustrate the intrinsic propensity of the evolved developmental process to produce phenotypic attractors that correspond to the two target phenotypes. e) *G* images that partially resemble *S2* in successively smaller patches (other elements of *G* are zero); adult phenotypes recall complete phenotype from each partial 'stimulus', up until the point where *G* is ambiguous. f) *G* images that partially resemble *S1* and partially resemble *S2* (in steps of 5% left to right); adult phenotypes still either fully resemble *S1* or fully resemble *S2*. g) As per (f) but with linear development (magnitudes of phenotypes rescaled for display). h) Two-dimensional projections of phenotype distributions: the position of each point is given by *D1*=Hamming distance to *S1*, *D2*=Hamming distance to *S2* (each phenotypic character is capped at +/- 1 for the purposes of this measure).

These phenotypes have N=1,764 genes (42x42 pixels) and accordingly, if the network were fully-connected, the number of interaction coefficients in B would be $N^2=3.1\cdot10^6$. This is too many to simulate evolution using low mutation rates in reasonable time and, in any case, exceeds the density of connections observed in real gene networks (e.g., Davidson 2010, Leclerc 2008). However, in Experiments 1 and 2 we find that for any one selective environment, after the signs of the elements in G have stabilized, the direction of selection on each interaction coefficient is constant (see *Analysis*). Under *directional* selection, the cumulative effect of a large number of small mutations is equivalent to the effect of a small number of large mutations when controlling for variance. In between changes in environments where selection is constant, we use this observation to reduce the number of mutation-selection cycles simulated (see Supplementary text (d)). In this example we also take the opportunity to demonstrate that it is not necessary for the gene network to be fully connected. Specifically, we evolve a sparsely connected network where each s_i has only 10 connections to other genes picked at random with equal probability.

As before, we can examine the effect that evolved interactions have on adult phenotypes by mutating G (Fig. 3.c-d). Notice that the output of the network is extremely robust to mutation in G (Fig. 3.c), far beyond what is easily recognisable as the target phenotype by visual inspection. In this sense, development can repair 'corrupted' embryonic phenotypes. When genetic information in G is completely destroyed (Fig. 3.d), development reproduces one or the other of the previously selected phenotypes (as before, the network also produces the complementary patterns; these are inverted to the 'positive' image for display). Fig. 3.e (also Supplementary text (j)) assesses adult phenotypes produced from embryonic phenotypes that *partially* resemble a previously selected phenotype; in this case, phenotypes that resemble the target patterns on only a small patch of the original image We see that the adult phenotypes reproduce the phenotype in its entirety. This effect is very robust up until the point that the patch becomes ambiguous at just a few pixels – then the phenotype produced may be either *S1* or *S2* with approximately equal probability. Thus as evolved phenotypic correlations become stronger, G need only specify a few traits of the embryonic phenotype in order for development to reproduce an adult phenotype in its entirety; Interestingly, this implies that selection on traits with direct effects (G) becomes less important to producing fit phenotypes. In Fig. 3.f the G is changed systematically from *S1* to *S2* (in steps of 5%) - we observe that the adult phenotypes change abruptly from *S1* to *S2* in response (Supplementary text (e)).

Linear versus non-linear developmental interactions

For the most part, this non-linear developmental mapping produces adult phenotypes *S1* or *S2* without 'mangling' the two patterns – i.e., it does not, in general, produce a low-fitness pattern that is the average of the two target patterns but produces either one high-fitness pattern fully or the other high-fitness pattern fully (Fig. 3.f, Supplementary text (e)). This recall of multiple distinct patterns cannot be achieved with a linear G-P mapping. In an experiment that is the same as above except that it uses a linear mapping, (i.e., $\sigma(x)=x$), the selective pressures on interaction coefficients are qualitatively similar and resultant interaction coefficients have the same signs and structure as before (i.e., they are still Hebbian); however, a linear developmental process is unable to reproduce either of the selected phenotypic patterns in a self-consistent manner and instead produces an average of the two targets (Fig. 3.g).

Fig. 3.h shows the distribution of left) target phenotypes, centre) the evolved adult phenotypes, and right) phenotypes produced by the linear G-P map. The phenotypic distribution from the non-linear developmental process is clearly bi-modal but the linear developmental system is essentially unimodal. In developmental terms the abrupt change in phenotype from one mode of this distribution to the other (e.g., Fig.3.f) represents genetic switching that produces a large specific change in phenotype in response to small changes in genotype (Supplementary text (e)). Crucially, although this developmental system is robust in the sense of only producing very specific combinations of characters, it is not canalised to produce only one phenotype; it is still capable of producing phenotypes that differ in many characters.

Logically, note that a unimodal distribution constitutes a memory that only remembers one thing – i.e., to the extent that such a distribution remembers the most recently selected phenotype all memories of previous selected phenotypes must be 'over-written' (or averaged together if development changes sufficiently slowly). In this sense, linear developmental processes are a degenerate sense of developmental memory since only a non-linear developmental process is capable of holding more than one memory.

The significance of non-linearity in the input-output transform function and its implications with respect to these behaviours is well-understood in the context of neural networks (both for feed-forward multi-layer networks and recurrent networks; Rumelhart & McClelland 1986, Hopfield 1982). Intuitively, a linear mapping is unable to 'break symmetry' or amplify differences that enable the developing phenotype to settle on either one target or the other, and accordingly this results in adult phenotypes that are a blend of the two target patterns. Relatedly, a *single-step* transform cannot produce two conflicting patterns without mangling them even if it is non-linear (this is related to the fact that a single-layer non-linear perceptron cannot represent XOR; Rumelhart & McClelland 1986). In contrast, with a non-linear recurrent mapping, positive feedback between developing characters can cause phenotypes to find a pattern that is consistent with one of the two selected patterns in a 'winner takes all' manner. In the context of cognitive/neural models a bi-modal or switch-like behaviour is referred to as 'recognition' or 'categorical perception' (Harnad 2005) – i.e., identifying which discrete class an input pattern belongs to (Fig.3.f).

Experiment 4) Structured variation: generalisation and modularity

If, as in Experiments 1-3, an evolved developmental process recalls phenotypes that have been selected in the past; does this mean that evolved development cannot produce anything new? Experiment 4 assesses the propensity of a developmental memory to generalise over a collection of phenotypic targets and produce novel phenotypes from the same class. The targets are drawn from a 'family' of related patterns (Parter, Kashtan & Alon 2008). This examines the ability to produce specific phenotypic patterns that have *not* been selected for in the past but that have features or underlying structure similar to those that *have* been selected in the past. Specifically, we anticipate that the new patterns will consist of novel combinations of phenotypic modules.

Experiment 4 uses the eight different target phenotypes (N=100) shown in Fig 4.a.. Each pattern has four subgroups of characters (the quartiles of each image) that vary in a simple modular fashion; i.e., characters within the same subgroup are strongly correlated whereas correlations between these groups are weak or absent (Lipson, Pollack, & Suh 2002, Wagner, Pavlicev & Cheverud 2007, Kashtan et al 2009, Lipson 2007, Watson et al. 2011b, Clune et al. 2013). In this example, each sub-pattern (a 'loop' or a 'stalk' in various orientations) appears in at least one of the target phenotypes. Again, the particular sub-patterns used for illustration are arbitrary but biologically these represent different phenotypic forms within a body plan – for example, petals and sepals (see also RNA loops and ladders used in Parter, Kashtan & Alon 2008). Here B is fully-connected (not sparse).

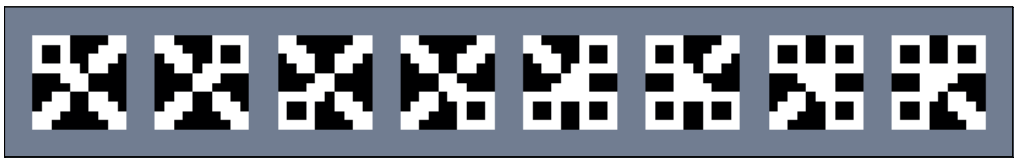

a) Eight target phenotypes.

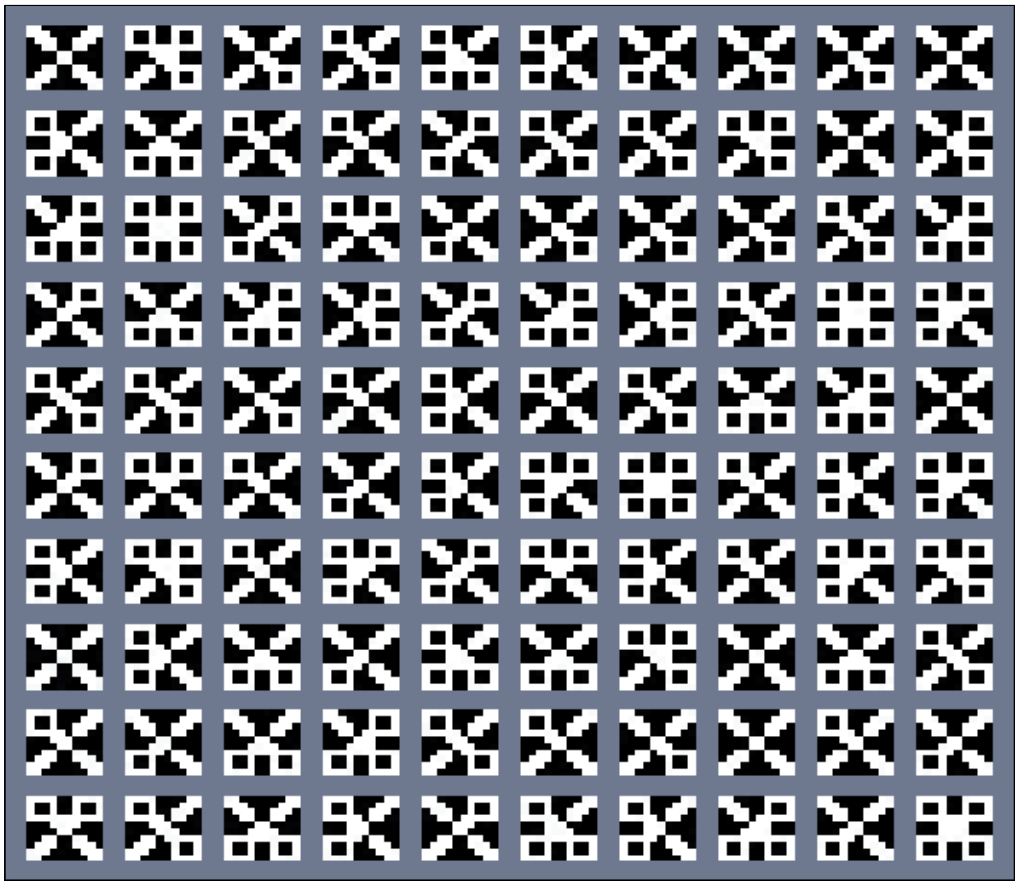

b) Bestiary of 100 adult phenotypes produced by evolved developmental network from random *G*.

c) 10 example adult phenotypes produced by linear developmental network.

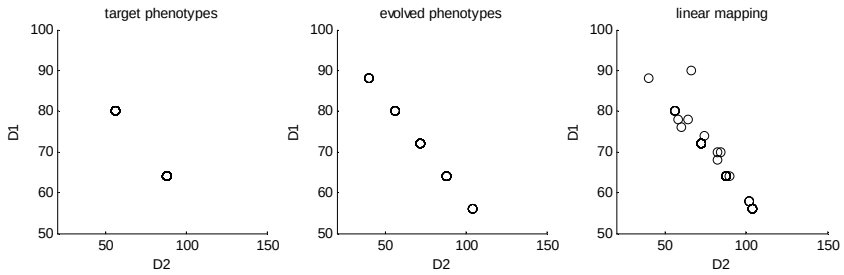

d) Phenotype distributions of targets and adults with non-linear and linear development.

Fig. 4. Experiment 4. Modularly-varying environment. a) Target phenotypes, *N*=100 genes (10x10 pixels) each, varying from one another in a modular fashion. b) Evolved adult phenotypes include the target phenotypes but also several generalisations. c) A linear developmental process produces many adult phenotypes that are incoherent. d) Two-dimensional projections of phenotype distributions: the position of each point is given by *D1*=Hamming distance from the 4-loop phenotype, *D2*= Hamming distance from the 0-loop phenotype (each phenotypic character is capped at +/- 1 for the purposes of this measure). The eight training patterns coincide on two points in this projection because these patterns differ by the addition/removal of exactly 1 loop or 3 loops.

Fig. 4.b examines the distribution of adult phenotypes produced from the evolved developmental interactions by randomising *G*, as before. We see that development produces each of the 8 training patterns, but it also produces simple generalisations of these phenotypes from new combinations of modules, e.g., phenotypes with four loops. Fig. 4.c. shows that again a linear developmental process is unable to produce this phenotypic distribution, instead producing phenotypes that are an approximate average of the training patterns. In Fig. 4.d., phenotypes are plotted as points in a two-dimensional projection of phenotype space. Phenotypes produced by the evolved developmental system include the training patterns but also specific other phenotypes; specifically, development generalises from target phenotypes that have 1 or 3 loops to produce phenotypes with 0, 1, 2, 3 or 4 loops. This shows that by producing new combinations of modules, developmental memory can generalise in both an interpolative and extrapolative manner from phenotypes that have been selected in the past. Supplementary text (f) discusses how this type of generalisation is (necessarily) equivalent to a 'failure' to restrict phenotypes to a set of training patterns accurately. Again the linear mapping (Fig. 4.d.right) also produces phenotypes that are intermediate within this range of possibilities.

Accordingly, we see that the evolved developmental process is not just reproducing previously selected phenotypic patterns, but internalising structural information about the set of target patterns – thus producing phenotypes that have *not been previously seen* but are in the same family of phenotypes (Parter, Kashtan & Alon 2008). In a different example the modularity in the evolved interaction matrix that enables this generalisation is easily observable (Supplementary text (g)).

Analysis and Discussion

Selection pressures on interaction coefficients are Hebbian

Here we show that when natural selection is sufficiently efficacious that phenotypic characters at least have the same sign as the direction of selection on those characters, beneficial changes in an interaction coefficient between those characters will follow Hebb's rule, i.e., $\Delta b_{ij} = r s_i s_j$, ($r$>0). In general, the selection coefficient of a mutation will be given by

$$\frac{w(\boldsymbol{P}^*+\Delta\boldsymbol{P}^*)-w(\boldsymbol{P}^*)}{w(\boldsymbol{P}^*)} = \frac{1+(\boldsymbol{P}^*+\Delta\boldsymbol{P}^*)\cdot\boldsymbol{S}}{1+\boldsymbol{P}^*\cdot\boldsymbol{S}} - 1 = \frac{\Delta\boldsymbol{P}^*\cdot\boldsymbol{S}}{\boldsymbol{1}+\boldsymbol{P}^*\cdot\boldsymbol{S}} ,$$

where ΔP^* is the vector of changes conferred on the adult phenotype. For a recurrent non-linear developmental process, a single mutation that alters an interaction coefficient by Δb_{ij} may affect many phenotypic characters. Instead, to build intuition, first imagine that development is represented by a single-step linear mapping (i.e., T=1, one iteration of Equation 1, and $\sigma(x) = x$); in this case the change in phenotypic character, p_x, due to a change in b_{ij}, is zero for all $x \neq i$. The phenotypic consequence of a mutational change Δb_{ij} is then

$$\Delta p_i = \left[p_i + \tau_1\left(\sum_{j=0}^{n}(b_{ij} + \Delta b_{ij})p_j\right) - \tau_2 p_i\right] - \left[p_i + \tau_1\left(\sum_{j=0}^{n} b_{ij}p_j\right) - \tau_2 p_i\right]$$

$$= \tau_1 \Delta b_{ij} p_j .$$

Thus the selective coefficient of the mutation is $\frac{\boldsymbol{\Delta p_i} s_i}{\boldsymbol{1}+\boldsymbol{p_i} s_i} = \frac{\tau_1 \Delta b_{ij} p_j s_i}{\boldsymbol{1}+\boldsymbol{p_i} s_i}$. Given that p_i and s_i have the same sign, this mutation is favoured if and only if $\tau_1 \Delta b_{ij} p_j s_i > 0$. Given $\tau_1 > 0$, a change Δb_{ij} is therefore beneficial if and only if it has the same sign as $s_i p_j$. Accordingly, for beneficial mutations we can write

the ratio of these quantities as, $\Delta b_{ij}/s_i p_j = r$, where r is some positive parameter. Thus all beneficial mutations satisfy the condition $\Delta b_{ij} = r\, s_i p_j$, ($r$>0). (We observe that for directional selection this agrees with the *Delta rule*, a simple supervised learning rule based on error minimisation where s_i is the desired direction of change in output i, and p_j is the input from node j). When the current value of the phenotypic character (p_j) agrees with the current direction of selection on that trait (s_j), all beneficial mutations satisfy the condition, $\Delta b_{ij} = r s_i s_j$, ($r$>0), as per Hebb's rule.

In the above experiments we do not assume that the sign of p_j matches the direction of selection, s_j. But when the evolution of G is significantly faster than the evolution of B, (as modelled here by assuming that heritable variation on interactions is lower than that on the direct effects), the embryonic traits quickly come to have the sign that is selected.

Although the algebra is far more complex when we relax our assumption that the adult phenotype is determined by a single iteration of a linear developmental function (because Δb_{ij} can then percolate into additional components of P^*), the simulations show that the same effect obtains using the multi-step non-linear developmental mapping, (i.e., T=10 and $\sigma(x)$=tanh(x)).

How surprised should we be that selection pressures on interaction coefficients are Hebbian?

Neural learning is usually conceived of as a mandated mechanism, applied with the intent of producing certain learning behaviours in neural networks; in contrast, the action of evolution by natural selection on relational traits is, of course, not mandated to follow these principles nor directed toward any predetermined function. However, the principle underlying the way that neural learning mechanisms adjust the strengths of neural connections is simply one of local incremental improvement, i.e., change each connection a little bit in the direction that makes the output more similar to the target. Accordingly, whenever heritable variation affects correlations, the direction of evolutionary change under natural selection, acting to improve the fitness of the phenotype, necessarily agrees with the principles of these neural learning mechanisms, as shown above.

Note that under natural selection, changes to interactions are selected because they make the current phenotype fitter – they cannot be selected because they produce a memory or for any other *future* consequence (including evolvability). But changing a correlation so that a given pattern is expressed more strongly or completely has the side-effect of making the network more likely to express that pattern again in future. In dynamical systems terms, these changes widen the basin of attraction for this pattern; i.e., increase the number of initial conditions (here, embryonic phenotypes) that lead to that configuration (Watson, Mills & Buckley 2010, Watson, Buckley & Mills 2010, Coolen 1991). This kind of learning thereby transforms correlations into causal interactions – that is, genes whose activation was originally coincident because of a correlated selection pressure come to have activations that are coincident because of internal regulatory interactions. Intuitively, the system finds things that were good together in the past and puts them together more often in the future; this simple principle is the essence of an associative memory, and from this principle all of the results we have shown follow.

Some of the general principles investigated here extend beyond either neural or ontogenetic interactions into other, quite different, domains (Watson, Mills & Buckley 2010, 2011, Davies et al 2010, Watson et al. 2014), including the tendency of natural selection on non-trophic ecological relationships to reinforce correlations in species densities and thereby produce 'ecosystem memory' (Lewis 2009, Watson et al. 2009, in prep, 2014).

Robustness/sensitivity to embryonic perturbation

Our experiments show that even for large *N* and sparse networks the distribution of phenotypes produced by the evolved developmental network is extremely robust to variation in the embryonic phenotype, $P(0)$, caused by, for example, mutations to *G*. Adult phenotypes would necessarily have equal robustness to any *environmental* variation that modified the embryonic phenotype in the same manner. If developing phenotypes are perturbed by environmental variation at later developmental time steps then there will be fewer remaining time steps in which development may correct these perturbations (a brief examination of this is shown in Supplementary text (h)).

In a multi-modal phenotypic distribution, robustness goes hand-in-hand with sensitivity – i.e., small variations applied to a developing embryo located at the saddle-point between two developmental attractors will necessarily produce large changes in the adult phenotype. In Fig.3.e. we artificially stimulated the embryonic phenotype to partially resemble one of the target phenotypes. But conceivably, an embryonic phenotype with even a small sensitivity to environment could similarly produce large changes in phenotype. This possibility makes an interesting link between phenotypic plasticity (West-Eberhard 2003) and developmental memory in that development may produce a complete 'preformed' phenotype in response to environmental cues rather than genetic variation.

Consequences of associative learning in the evolution of development

In this paper we have demonstrated a formal equivalence between the direction of selection on phenotypic correlations and associative learning mechanisms. In the context of neural network research and connectionist models of memory and learning, simple associative learning with the ability to produce an associative memory, to store and recall multiple patterns, categorise patterns from partial or corrupted stimuli, and produce generalised patterns from a set of structurally similar training patterns has been well studied (e.g., Hopfield 1982, Rumelhart et al. 1986, O'Reilly & Munakata 2000). The insight that the selective pressures on developmental correlations are equivalent to associative learning thus provides the opportunity to utilise well-established theoretical and conceptual frameworks from associative learning theory to identify organisational principles involved in the evolution of development; e.g., to understand the minimal conditions for a developmental memory capable of the behaviours illustrated above. From this it follows that evolved developmental processes can exhibit learning and memory with the same affordances and limitations as the manner in which associative learning mechanisms cause a neural network to form a memory of a set of training patterns. This provides a specific example of the more general formal connection between evolution and learning (Valiant 2009). Accordingly, the idea that gene networks can act in a manner analogous to connectionist models of cognitive behaviours is more than merely a metaphor, and helps us make sense of how biological networks evolve adaptive complexity (Stewart 1997, Sansom 2011).

This paper has focussed on how evolution affects development. But logically, any predisposition of a developmental process to produce particular phenotypes rather than others may affect the speed and direction of subsequent evolution. In particular, when an evolved developmental process is primed to produce particular phenotypes that have been selected for in the past (as shown here) this can facilitate a population in evolving such phenotypes (or similar phenotypes) should they be selected for in the future (Stewart 1997, Kashtan, Noor & Alon 2007, Parter, Kashtan & Alon 2008). Exactly how the evolved properties of development affect such evolvability (Wagner & Altenberg 1996, Kirchner & Gerhart 1998, Wagner, Pavlicev & Cheverud 2007, Wagner & Laubichler 2004, Pavlicev & Wagner 2012, Draghi & Wagner 2008/2009, Hendrikse, Parsons & Hallgrímsson 2007, Izquierdo & Fernando 2008, Laland, et al. 2011) will be analysed in detail separately. The challenge in this area is to understand how historical contingency affects future evolutionary adaptation, and in particular how increasing genetic constraint can

improve adaptability (Wimsatt 2001). Making a formal connection between the evolution of development and the principles of memory and learning, in this manner, provides access to a theoretical framework for tackling these questions (Valiant 2009). In particular, the principles of generalisation, such as the tension between short-term performance gains and long-term performance losses due to over-fitting, may be useful for understanding the conditions for and limitations of evolvability (Supplementary text (f) & (i)).

The fact that natural selection can alter the distribution of phenotypic variation, and reflexively, that the distribution of phenotypic variation can alter the selective pressures on subsequent evolutionary change, is an example of 'reciprocal causation' in evolution (Laland, et al. 2011). Conceiving evolution as a learning process, rather than a fixed trial and error process, helps to explain how evolution can alter its own mechanisms in this reciprocal sense (Watson et al. 2014). Specifically, the equivalence of the selective pressures on ontogenetic interactions with associative learning mechanisms demonstrated here illustrates how evolution can respond to future selection in a manner that is 'informed' by past selection in exactly the same sense that cognitive learning systems are informed by past experience.

Acknowledgements: Thanks to Adam Davies, Chris Buckley, Rene Doursat, Nick Barton, Loizos Kounios, Tom Ezard and Jeff Clune for discussion, and Kostas Kouvaris for technical contributions. This work in part supported by National Science Foundation Emerging Frontiers Award 1038657 to DMW.

Supplementary text (a-j) is provided in on-line supplementary materials.

References

Ackley, D.H., Hinton, G.E., Sejnowski, T.J. (1985) A Learning Algorithm for Boltzmann Machines, *Cognitive Science*, 9: 147-169.

Brigandt, I. (2007) Typology now: homology and developmental constraints explain evolvability. *Biol Philos (2007) 22:709–725.*

Cheverud, J. M., Ehrich, T. H., Vaughn, T. T., Koreishi, S. F., Linsey, R. B., & Pletscher, L. S. (2004). Pleiotropic effects on mandibular morphology II: differential epistasis and genetic variation in morphological integration. *Journal of Experimental Zoology Part B: Molecular and Developmental Evolution*,302(5), 424-435.

Chevillon, C., Bourget, D., Rousset, F., Pasteur, N. & Raymon, M. (1997) Pleiotropy of adaptive changes in populations: comparisons among insecticide resistance genes in Culex pipiens. *Genet Res*, 195-204.

Clune, J., Mouret, J.B., & Lipson, H. (2013). The evolutionary origins of modularity. *Proceedings Royal Society B*, 280(1755).

Coolen, A.C.C. (1991) On the Relation Between Stability Parameters and Sizes of Domains of Attraction in Attractor Neural Networks. *Europhys. Lett.* 16. 73.

Crombach. A., Hogeweg, P. (2008) Evolution of Evolvability in Gene Regulatory Networks. *PLoS Comput Biol* 4(7).

Davidson, E. H. (2010). Emerging properties of animal gene regulatory networks. *Nature*, 468(7326), 911-920.

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. *Artificial Life* 17 (3), 167-181.

Delph, L. F., Steven, J. C., Anderson, I. A., Herlihy, C. R., & Brodie III, E. D. (2011). Elimination of a Genetic Correlation Between the Sexes Via Artificial Correlational Selection. *Evolution*, 65(10), 2872-2880.

Draghi, J. & Wagner G.P. (2008) Evolution of evolvability in a developmental model. *Evolution* 62, 301–315.

Draghi, J. & Wagner, G. P. (2009). The evolutionary dynamics of evolvability in a gene network model. *Journal of Evolutionary Biology*, 22(3):599-611.

Draghi, J. A., Parsons, T. L., Wagner, G. P., & Plotkin, J. B. (2010). Mutational robustness can facilitate adaptation. *Nature*, 463(7279):353-355.

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.

Gillespie, J. H. (1984) Molecular evolution over the mutational landscape. *Evolution* 38, 1116-1129.

Guillaume, F., & Otto, S.P. (2012). Gene Functional Trade-offs and the Evolution of Pleiotropy. *Genetics*.

Harnad, S. (2005) "To Cognize is to Categorize: Cognition is Categorization". In C Lefebvre & H. Cohen. *Handbook of Categorization in Cognitive Science.* New York: Elsevier Press

Hebb, D.O. (1949) *The organization of behaviour*. New York: Wiley.

Hendrikse, J.L., Parsons, T.E., & Hallgrímsson, B. (2007). Evolvability as the proper focus of evolutionary developmental biology. *Evol Dev.*, 9(4):393-401

Hopfield, J.J. (1982) Neural networks and physical systems with emergent collective computational abilities, *PNAS USA*, 79 (8) 2554-2558.

Huang, S., Eichler, G., Bar-Yam, Y., & Ingber, D. E. (2005). Cell fates as high-dimensional attractor states of a complex gene regulatory network. *Physical review letters*, 94(12), 128701.

Izquierdo, E.J. & Fernando, C. (2008) The Evolution of Evolvability in Gene Transcription Networks. *Proceedings of Alife XI*.

Jones, A. G., Arnold, S. J. & Bürger, R. (2007) The mutation matrix and the evolution of evolvability. *Evolution* 61, 727–745.

Kashtan, N. & Alon, U. (2005) Spontaneous evolution of modularity and network motifs, *PNAS USA*, 102 (39) 13773-13778.

Kashtan, N., Mayo, A., Kalisky, T. & Alon, U. (2009) An Analytically Solvable Model for Rapid Evolution of Modular Structure. *PLoS Comput Biol* 5(4)

Kashtan, N., Noor, E. & Alon, U. (2007) Varying Environments Can Speed up Evolution. *Proceedings of the National Academy of Sciences of the United States of America* , Vol. 104, No. 34 (Aug. 21, 2007), pp. 13711-13716.

Kauffman, S.A. (1993) *The Origins of Order* (Oxford University Press, New York).

Kim, H.S., Huh, J., & Fay, J. C. (2009). Dissecting the pleiotropic consequences of a quantitative trait nucleotide. *FEMS yeast research*, 9(5), 713-722.

Kirchner, M. & Gerhart, J. (1998). Evolvability. *Proc. Natl. Acad. Sci. USA*, 95:8420–8427.

Laland, K. N., Sterelny, K., Odling-Smee, J., Hoppitt, W., & Uller, T. (2011). Cause and effect in biology revisited: is Mayr's proximate-ultimate dichotomy still useful? *Science*, 334(6062), 1512-1516

Lande, R. & Arnold, S. J. (1983) The measurement of selection on correlated characters. Evolution 37, 1210–1226.

Leamy, L. J., Pomp, D., & Lightfoot, J. T. (2009). Genetic variation in the pleiotropic association between physical activity and body weight in mice. *Genet Sel Evol*, 41, 41.

Leclerc, R. D. (2008). Survival of the sparsest: robust gene networks are parsimonious. *Molecular systems biology*, 4(1).

Lenski, R. E. (1988a). Experimental studies of pleiotropy and epistasis in *Escherichia coli*. I. Variation in competitive fitness among mutants resistant to virus T4. *Evolution*, 425-432.

Lenski, R. E. (1988b). Experimental studies of pleiotropy and epistasis in *Escherichia coli*. II. Compensation for maldaptive effects associated with resistance to virus T4. *Evolution*, 433-440.

Lewis, M. (2009) *An Investigation Into The Evolution Of Relationships Between Species In An Ecosystem*. MSc Dissertation, ECS, University of Southampton.

Lipson H., (2007) Principles of modularity, regularity, and hierarchy for scalable systems, Journal of Biological Physics and Chemistry, vol. 7, no. 4, pp. 125-128.

Lipson, H., Pollack, J.B., & Suh, N.P. (2002) On the Origin of Modular Variation. *Evolution*. Volume 56, Issue 8, pages 1549–1556, August 2002.

McGregor, S., Vasas, V., Husbands, P., Fernando, C. (2012) Evolution of Associative Learning in Chemical Networks, *PLoS Computation Biology*, 8:11 (e1002739)

Mills, R. (2010) *How Micro-Evolution Can Guide Macro-Evolution: Multi-Scale Search via Evolved Modular Variation.* PhD thesis, ECS, University of Southampton.

Mills, R, Jansen, T. & Watson, R.A. (submitted) *Transforming Evolutionary Search Into Higher-Level Evolutionary Search by Capturing Problem Structure*.

Mills, R., & Watson, R.A. (2009) Symbiosis Enables the Evolution of Rare Complexes in Structured Environments. *Proceedings of 10th European Conference on Artificial Life (ECAL 2009)*, 110-117.

O'Reilly, R. C. & Munakata, Y. (2000) *Computational Explorations in Cognitive Neuro-science: Understanding the Mind by Simulating the Brain* (MIT Press, Cambridge, MA).

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.

Pavlicev, M. & Wagner, G.P. (2012) A model of developmental evolution: selection, pleiotropy and compensation. *Trends in Ecology & Evolution*, Volume 27, Issue 6, 316-322

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* 62 (1), 199-213

Pavlicev, M., JM Cheverud, GP Wagner. (2011) Evolution of adaptive phenotypic variation patterns by direct selection for evolvability *Proceedings of the Royal Society B: Biological Sciences* 278 (1713), 1903-1912.

Raff, R.A. (2000). Evo-devo: the evolution of a new discipline. *Nature Rev. Genet*. 1: 74–79.

Riedl, R. (1978). *Order in Living Organisms*, Wiley.

Rumelhart, D.E., James L. McClelland and the PDP Research Group (1986) *Parallel Distributed Processing.* Vols. I and II. MIT Press.

Sansom, R. (2011). *Ingenious Genes: How Gene Regulation Networks Evolve to Control Development*. MIT Press.

Schluter, D. (1996). Adaptive radiation along genetic lines of least resistance. *Evolution* 50, 1766–1774.

Sniegowski, P.D., Murphy, H.A. (2006). Evolvability. *Current Biology* 16:R831-R834.

Stewart, J. (1997). The evolution of genetic cognition. *Journal of Social and Evolutionary Systems*, 20(1), 53-73.

Toussaint, M., & von Seelen, W. (2007) Complex adaptation and system structure, *BioSystems* 90: 769–782

Valiant, L. (2009). Evolvability, *J. Assoc. Computing Machinery*, 56:1 (2009) 3:1 - 3:21.

Vohradsky, J. (2001a). Neural model of the genetic network. *The Journal of Biological Chemistry*, 276:36168-36173.

Vohradsky, J. (2001b). Neural network model of gene expression. *The FASEB journal*, 15(3):846-854.

Waddington, C. H., & Robertson, E. 1966. Selection for developmental canalisation. *Genet. Res.* 7: 303–312.

Wagner, G. P. & Laubichler, M.D. (2004) Rupert Riedl and the re-synthesis of evolutionary and developmental biology: body plans and evolvability. *J Exp Zool B: Mol Dev Evol.* 302B:92-102

Wagner, G. P. (1989). Multivariate mutation-selection balance with constrained pleiotropic effects. Genetics 122, 223–234

Wagner G.P. & L. Altenberg, L. (1996). Complex adaptations and the evolution of evolvability. *Evolution* 50:967–976.

Wagner, G.P., Pavlicev, M., & Cheverud, J.M. (2007) The road to modularity. *Nature Reviews Genetics* 8: 921-931.

Watson, R.A., Palmius, N., Mills, R., Powers, S.T., & Penn, A.S. (2009) Can Selfish Symbioses Effect Higher-level Selection? *Proceedings of 10th European Conference on Artificial Life (ECAL 2009)*. 27-36.

Watson, R.A., Buckley, C. L., Mills, R., Davies, A.P. (2010a) Associative memory in gene regulation networks, *ALIFE XII*, 194-202.

Watson, R.A., Mills, R., Buckley, C. L., Penn, A., Davies, A., Noble, J. and Bullock, S. (2010) Adaptation without natural selection. in *Artificial Life XII: Proceedings of the Twelfth International Conference on the Synthesis and Simulation of Living Systems*, MIT Press, 80-81

Watson, R.A., Szathmary, E., Power, D., Mills, R., Powers, S.T. and Czapp, B. (in prep) *The Evolution of Ecological Relationships and Principles of Distributed Learning in Ecosystems*. Manuscript in preparation.

Watson, R.A., Buckley, C. L. and Mills, R. (2010) Optimisation in 'Self-modelling' Complex Adaptive Systems. *Complexity*. 16, (5), 17-26.

Watson, R.A., Mills, R. & Buckley, C. L. (2010) Global Adaptation in Networks of Selfish Components: Emergent Associative Memory at the System Scale. *Artificial Life*, 17, (3), 147-166.

Watson R.A., Mills, R., Buckley, C.L. (2011) Transformations in the Scale of Behaviour and the Global Optimisation of Constraints in Adaptive Networks. *Adaptive Behavior* 19, (4), 227-249.

Watson, R.A., Weinreich, D.M. & Wakeley, J. (2011) Genome Structure and the Benefit of Sex. *Evolution* 65, (2), 523-536.

Watson, R.A., Jackson, A., Palmius, N., Mills, R.M. & Powers, S.T. (submitted-a) *The Evolution of Symbiotic Partnerships, the Adaptive Capabilities of Ecosystems and the Algorithmic Principles underlying the Major Evolutionary Transitions*

Watson, R.A., R. Mills, R., C.L. Buckley, K. Kouvaris, A. Jackson, S.T. Powers, C. Cox, S. Tudge, P. Ryan, A. Davies (2014) *Evolutionary Connectionism: How the principles of associative learning operate in the evolution of developmental interactions, the evolution of ecological relationships and the evolution of new levels of individuality*. Technical Report, ECS, University of Southampton, Southampton, U.K.

Wessels, L., van Someren, E., and Reinders, M. (2001). A comparison of genetic network models. *In Pacific Symposium on Biocomputing*, number 6, pages 508-519.

West-Eberhard, M-J. (2003). *Developmental plasticity and evolution*. Oxford University Press, New York.

Wimsatt, W.C. (2001). Generative entrenchment and the developmental systems approach to evolutionary processes. *Cycles of contingency: Developmental systems and evolution*, 219-37.

Young, N.M., & Hallgrímsson, B. (2005). Serial homology and the evolution of mammalian limb covariation structure. *Evolution*, 59(12), 2691-2704.

Young, N.M., Wagner, G.P., & Hallgrímsson, B. (2010). Development and the evolvability of human limbs. *Proceedings of the National Academy of Sciences*,107(8), 3400-3405.

Supplementary Text

To accompany manuscript:

The evolution of phenotypic correlations and "developmental memory"

Richard A. Watson, Günter P. Wagner, Mihaela Pavlicev, Daniel M. Weinreich, Rob Mills

(a) Mean and relational effects of interaction terms

In the developmental system we model, when all interaction terms in *B* are zero, changes to *G* have an effect on the mean value of the phenotypic traits (including their sign) but not their correlation; hence we refer to *G* as the 'direct' effects of the genotype on phenotypic characters. Conversely, for a distribution of *G* elements each with a mean of zero, the effect of *B* is to alter the correlation of phenotypic traits but not their mean. However, as expected for general epistatic models (and in contrast to statistical models that keep interaction terms and mean effects separate), whenever trait values are non-zero a change to an interaction term can have both a 'mean effect' as well as a 'relational effect'. This is important because interaction coefficients that have no mean effect on phenotype can only be selected because of their effect on the shape of the distribution of phenotypes they produce (Pavlicev, Cheverud & Wagner 2011), but interaction coefficients that also have a mean effect on phenotypes can be selected because of their mean effect but nonetheless have a (systematically related) effect on the shape of the distribution of phenotypes.

(b) Hill-climbing model of selection

The evolution of the population is modelled by assuming the introduction of a mutant genotype, G' and/or B', being small mutations of $\bar{G}$ and/or $\bar{B}$, (in single copy number in a large population). This genotypic mutation uniquely determines a mutant phenotype, $P^{*\prime}$. In general, the probability of this mutant fixing in the population is proportional to the selective advantage of that mutation. But, here we aim to provide a phenomenological model of the evolution of development and as such it is the direction of selection, rather than quantified rates, that are important (rates will be sensitive to parameters such as magnitude of selection coefficients, mutation probabilities and the number of genetic loci contributing to a quantitative trait, etc. which will be case specific). Accordingly, it is sufficient for our purposes to assume a 'hill-climbing' model of selection (e.g., Kashtan et al. 2009), i.e., that the selection coefficient is sufficiently large that beneficial mutations fix and deleterious mutations do not. In general, mutation-rate-limited evolution can have qualitatively different dynamic properties to selection-intensity-limited evolution; but in all our experiments we find that the direction of selection using the former agrees with analysis using the latter (Pavlicev, Cheverud & Wagner (2011), and the stochastic nature of selection that would occur with small selection coefficients does not add to the clarity of exposition. Thus, if the mutant genotype is beneficial (i.e., $fitness(develop(G', B', T)) > fitness(develop(\bar{G}, \bar{B}, T))$, this becomes the mean genotype of the population ($\bar{G}(t+1)= G'$, $\bar{B}(t+1)=B'$), otherwise the loss of the mutant from the population leaves the mean genotype of the population unchanged. Only the fitness rankings of phenotypes (not the magnitudes of the fitnesses) are relevant for a hill-climbing selection model.

(c) Effect of evolved interaction terms on the shape of phenotypic distributions

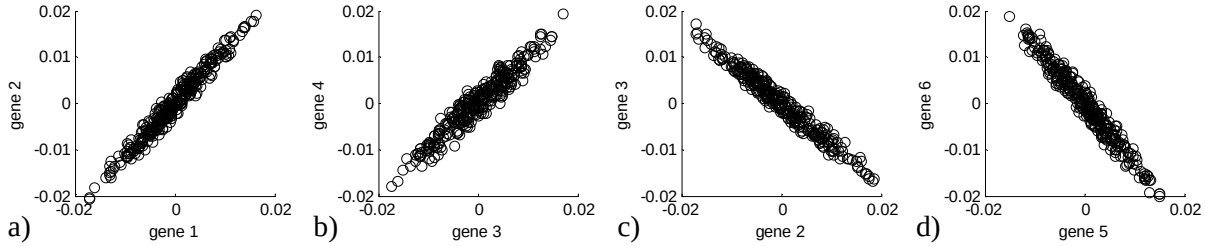

Fig. S1: The effect of evolved regulatory interactions on the correlation of expression. Experiment 1, after 10,000 generations, target, S = ++---+-+; Effect on correlation for a pair of genes that are selected together s_{12} = ++ (a), s_{34} = -- (b), or selected contrariwise, s_{23} = +- (c), s_{56} = -+ (d). Each point in each distribution shows the expression levels of two genes from a phenotype developed using the evolved regulatory interactions. Each phenotype is developed from a random G drawn from a uniform distribution.

Figure S1 confirms that the correlation of characters in the phenotypic distribution produced by the evolved developmental interactions agrees with the direction of selection experienced on those two characters in the selective environment.

(d) The cumulative effect of many small mutations under directional selection.

Under *directional* selection, the cumulative effect of a large number of small mutations is equivalent to the effect of a small number of large mutations when controlling for variance. When the time for G to stabilise after a change in environment is small compared to the number of generations after it stabilises before the next change in environment, we can therefore model the effect of natural selection in any one selective environment as follows. For each element of B (in random order without replacement) we test for both positive and negative directional selection by applying hill-climbing selection on two mutations, one drawn from a distribution with mean q, the other with mean $-q$, each with standard deviation σ. Under directional selection, at most one of these mutations will be retained by selection. Experiments 3 and 4 use q=0.02 and q=0.005, respectively. In both cases we use a standard deviation σ=0.01q.

A standard deviation of 1% of the mean simulates the cumulative effect of 10^4 discrete mutations (each with equal probability of being positive or negative) occurring before each change in environment (coming from $\sigma = \sqrt{np(1-p)},\ (p = 0.5)$). Naturally, there is a trade-off in these experiments between how many mutations occur before each change in environment and how many changes in environment are simulated in total. If the latter is increased, the former can be decreased, but this requires simulating more mutation-selection cycles. Here we find that 40 changes in environment is sufficient for Experiment 3 whereas Experiment 4 uses lower mutation magnitudes and 800 changes in environment (on average, 100 presentations of each training pattern). This low mutation rate enables the evolution of the 16 phenotypic attractors without 'over-fitting' to the 8 target patterns (see section (f)).

(e) Developmental switching – abrupt change in phenotype from linear change in genotype

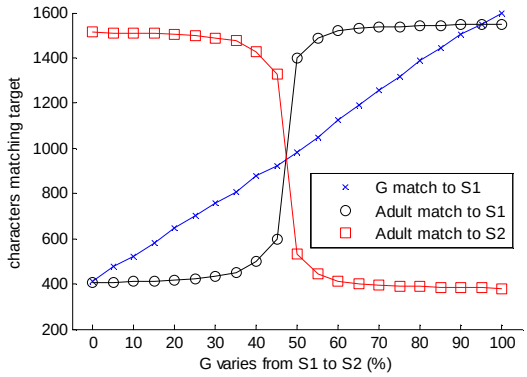

Fig. S2: Quantified detail of Fig.3.f. – match of *G* and adult phenotypes to targets *S1* and *S2*.

In an experiment where *G* varies systematically from *S1* to *S2* in steps of 5% (Fig. 3.f), note that as *G* changes slowly from *S1* to *S2* the adult phenotype switches abruptly from *S1* to *S2*.

(f) Generalisation and evolvability

Any kind of canalisation seems to oppose the possibility that a developmental representation can be primed to produce phenotypic patterns that are genuinely new. However evolved correlations – or more exactly, the restrained application of evolved correlations – can facilitate evolutionary novelty in a quantifiable sense. Specifically, we can use knowledge of neural network learning to understand the affordances and limitations of enhanced evolvability to produce outputs that are new. Here we equate the selective environment that the evolving GRN has already been exposed to with the *training set* of the learning neural network. High fitness in these selective environments corresponds to high performance on (the current element of) the training set. Whereas, the ability to evolve high fitness phenotypes in *new* selective environments is analogous to the ability of a well-trained neural network to *generalise* to an unseen *test set*. For example, in Experiment 4, the 'training set', contains 8 patterns but we observe that the network is predisposed to create 16 phenotypic patterns, not just these 8.

This type of generalisation is not to be taken for granted, however, even when the training set contains the structural patterns that are representative of the general class. The problem is that the training set, being only a subset of the general class, necessarily also contains other structural patterns that you do not want the system to learn if it is to provide good generalisation. A learner may in some cases learn the specific patterns of the training set and thereby fail to generalise to the test set. This is called *over-fitting* (Rumelhart et al. 1986). In general, the difficult part of machine learning is not the task of getting high performance on the training set but the task of avoiding over-fitting to get good generalisation on the test set. The issue is exacerbated if the model space is particularly 'expressive' or high-dimensional (as is necessary for complex learning tasks). In this case the model space will be highly *under-constrained* – meaning that there are many ways to represent the training set perfectly. Unfortunately, most of these solutions do not generalise well; so the task of finding general solutions is all the more difficult when the model space is high-dimensional.

Note that any kind of generalisation is equivalent to a 'failure' of a memory to represent *only* the training set. One solution to over-fitting is therefore to limit the complexity of the model space such that the learning system can only attain high fitness on the training set by finding general solutions; in other words, such that it does not have the capacity to learn the patterns by rote. In neural network models, model complexity is sometimes controlled by adding a penalty for each neural connection that is used (O'Reilly & Munakata 2000). Similarly, Clune, Mouret & Lipson (2013) add a cost to connections and observed improved evolvability, and Kashtan, Noor & Alon (2007) simply limit the number of logic gates

a solution can use. In the current work, the training set of Experiment 4 has no (pairwise) correlations between modules and thus there is nothing for the GRN to learn. Or more exactly, there is no structure that the GRN is *capable* of learning. [2]

In the context of neural learning, when a memory is formed that was not in the training set this is sometimes referred to as a "spurious memory". But spurious memories can also be interpreted as a simple form of generalisation – producing a pattern that contains features of several different training patterns (Fontanari 1990, Watson, Buckley & Mills 2009, Watson, Mills & Buckley 2010, Jang, Kim & Lee 1992). In some cases, this generalisation ability has the effect of identifying meso-scale features that are common to many patterns and putting them together in new combinations. This can create a genuinely novel pattern that is substantially different from any of the training patterns, and yet contains features that resemble sub-patterns observed in multiple training patterns (Watson, Mills & Buckley 2010), as illustrated in Experiment 4. It is no coincidence that this kind of generalisation involves modules. Consider two traits in different modules; If generalisation is to be permitted then the GRN must not enforce a correlation between these two traits. In fact, for *any* two traits in different modules, the GRN knows *nothing* about what combinations of traits are fit. Accordingly, the only way for the GRN to know *anything* about past selective environments (whilst still allowing some traits to vary independently) is for it to learn correlations in approximately disjoint subsets. Learning correlations within modules usefully restricts the phenotypes that are produced, but learning correlations between modules opposes generalisation and novelty.

Note that learning the correlations within modules improves performance on the training set, but the only advantage of *not* learning the correlations between modules, allowing modules to vary independently of one another, is for improving generalisation and performance on the test set. This suggests that the selective pressures on learning what things go together are quite different from the selective pressure to learn what things to keep separate ('parcelation' rather than 'integration'; Wagner & Altenberg, 1996). This perspective indicates that the challenge for understanding the evolution of evolvability is closely related to the difficulty of learning general models from data (Valiant 2009) – in particular, overcoming the problem of over-fitted solutions that give high-performance on the training set but fail to generalise to unseen test cases. And more specifically, the capability of expressive of G-P mappings to represent complex structural patterns is balanced by their tendency to over-fit to the past and thus fail to generalise to future selective environments.

Nonetheless, the equivalence between developmental memory and associative learning demonstrated in the current paper indicates that the evolution of a G-P mapping can provide substantive adaptive advantages in quite general scenarios. For example, the capability of recurrent neural networks to find locally optimal solutions to constraint optimisation problems (Hopfield & Tank 1986) is well known and associative learning can improve this ability (Watson, Buckley & Mills 2010, Watson, Mills & Buckley 2011). Specifically, associative learning can improve the ability of a network to find high-quality solutions even when the problem has only 'accidental' structure rather than an explicit modular decomposition (Watson, Buckley & Mills 2011). But the advantage is much more pronounced when there is modular structure that can be exploited (especially if the G-P map changes the variational neighbourhood, enabling large but controlled variation in phenotype space) (Watson, Mills & Buckley 2011, Mills 2010, Mills & Watson 2009, Mills, Jansen & Watson submitted). The potential for evolution by natural selection, given developmental processes with suitable heritable variation, to exhibit such

[2] In Experiment 4, the number of loops in the training patterns represent *odd parity* – i.e., the training samples have 1 or 3 modules. Odd parity is a generalisation of XOR (logical Exclusive OR) to multiple inputs that famously cannot be solved by local pairwise interactions. Since the developmental network can represent pairwise correlations, but not higher-order correlations, it cannot learn this set of training patterns without also producing patterns with 0, 2 and 4 loops.

adaptive capabilities suggests that the evolution of evolvability may alter the adaptive capabilities of natural selection in a substantial manner; and, perhaps more importantly, the application of a machine learning framework enables us to quantify and understand the affordances and limitations of this ability (Watson et al. 2010a, Watson et al. 2014).

(g) Modularly varying environments produces modular structure in the interaction network

In Experiment 4 in the main text the two patterns in each module are arbitrary, not the complement of one another, and different patterns are used in each module (e.g., different rotations of loops). The characters that correspond to a block are also not continuous in the linear genome *G* (there is nothing about the experiment that is sensitive to this ordering of the phenotypic characters[3]). This serves to illustrate that the result does not require special symmetries. Here we provide an additional experiment, analogous to Experiment 4, but with more transparent modularity that can be easily verified by inspecting *B*. We studied 8 target phenotypes, each of which is composed of 16 characters, divided into four groups of four (Fig. S3.a). The four characters within each group only appear in the set of target patterns as either ---- or ++++. Four of these sub-patterns are combined to create a complete training pattern of 16 characters such that each target pattern contains one module of the "++++" type and the remainder are of the "----" type or vice versa.

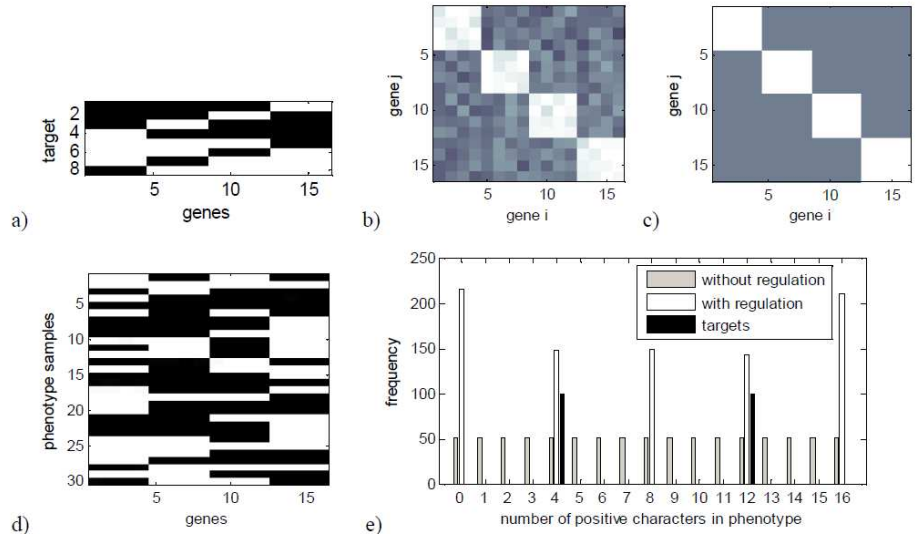

Fig. S3: Interaction coefficients evolved in modularly varying selective environment ($9.6 \cdot 10^6$ generations) and resultant phenotypes. a) The set of modularly varying target phenotypes. b) Evolved interaction coefficients. c) Hebbian interaction matrix summed over the 8 patterns. Evolved within-module interactions are positive and between-module interactions are approximately zero, exactly as expected under Hebbian learning. d) 30 example adult phenotypes developed from uniformly random *G*. Note that these include patterns that are not in the target set. e) Distributions of phenotype patterns; 'Targets' have either 4 or 12 positive traits. 'Without regulation' shows the frequency positive traits in *G*. For this figure only, *G* is are generated to be uniformly distributed in number of positive characters [0,16]; 'With regulation' shows the frequency of '+' traits in the adult phenotype patterns developed with

[3] The ability of development to produce novel phenotypes via new combinations of developmental modules is closely analogous to the ability of sexual recombination to produce novel genotypes via new combinations of genetic building blocks (Watson, Weinreich & Wakeley 2011), and the test case used here is structurally equivalent. However, the success of sexual recombination requires that genetic linkage corresponds with epistatic dependencies, whereas here useful modules can be identified and reproduced in new combinations regardless of the ordering of traits in the target phenotypes.

the evolved interaction coefficients. Note that the phenotypes produced include the target patterns (with one or three blocks of positive traits) but also all other combinations of modules (with zero, two and four blocks of positive traits).

Figure S3.b-c shows that the evolved correlations clearly identify the strong correlations amongst the genes within each module as predicted by Hebb's rule. We then examine the distribution of adult phenotypes produced from these evolved correlations (Fig. S3.d-e). We see that development reproduces only well-formed modules (i.e., the traits within each module match either '++++' or '----'), but that in addition to the target phenotypes, the adult phenotypes also exhibit other patterns produced by new combinations of those modules. In fact, all 16 possible combinations of the four modules are exhibited in the distribution of adult phenotypes (not just the 8 in the training set). In this simple example we can see the modularity in *B* (Fig. S3.b) (representing the fact that there are strong correlations between traits in the same module and no evolved correlations between traits in different modules) that thus allows novel combinations of modules to be produced by development.

(h) Brief investigation of robustness to environmental noise

Each row is a sample of 10 phenotypes (as per Fig. 3.d) but each row shows the effect of environmental perturbation (in this extreme case, a full randomisation of the values of *P*) at progressively later time steps of development. For example, when environment affects *P* only 3 time steps from the end of development, the adult phenotype (at time step 10) is unable to recover either of the target phenotypes accurately.

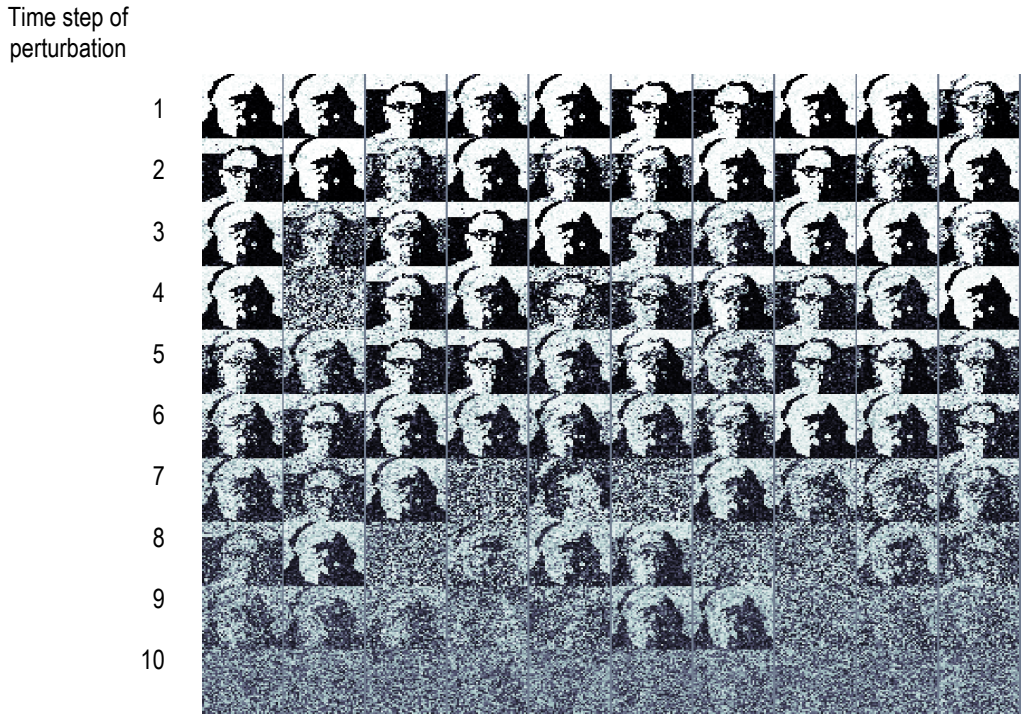

Fig. S4: Example phenotypes at developmental time step 10, after an environmental perturbation (randomisation) to the developing gene expression pattern at successively later developmental time steps. Eeach image is an independent simulation, 10 examples for each perturbation period.

However, note that since there are only two phenotypic attractors in this developmental system, and any *G* leads to either one or the other (Fig. 3.d), if development were allowed to continue for additional time

steps after an environmental perturbation, in every case it will eventually equilibrate at one or other of the two targets as before. In these experiments we are interpreting environmental perturbation simplistically; a) we are investigating whether an adult phenotype can recover from an environmental perturbation not whether a phenotype can resist environmental perturbation in the first place, b) here we are also not investigating what happens when environmental perturbation of the phenotype is partial.

(i) Robustness without reducing phenotypic variability

An evolved developmental process with appropriate flexibility may enable the recovery of high-fitness phenotypes quickly and reliably after a change in environment. The converse is also true, however; the ability of a population to evolve a particular phenotype may be retarded by a developmental process that does not produce appropriate phenotypic variation. This has some relevance to the apparent tension between evolvability and robustness. That is, the conservative and myopic quality of natural selection favours developmental canalisation and robustness which seems intrinsically opposed to an increase in phenotypic variability or evolvability. Some view these notions as two-sides of the same coin (Kirchner & Gerhart 1998, Draghi et al. 2010, Brigandt 2007); i.e., a predisposition to evolve some phenotypes more readily goes hand in hand with a decrease in the propensity to produce other phenotypes. Indeed, the regulatory connections that evolve in our model do not evolve because they increase the speed of future evolution, although they do, but rather because, in the immediate term, they increase the strength with which the current phenotype is expressed (see discussion of training set and test set in section f).

A multi-dimensional concept of variability, including correlated variability, provides a more sophisticated appreciation of the manner by which canalisation might oppose or facilitate evolvability. For example, if robustness is increased by canalisation applied uniformly in all dimensions then that would oppose evolvability. But there are (at least) two other possibilities. i) An increase in robustness might occur by reducing variability in some traits whilst leaving others free to vary; focussing the remaining variability on more useful dimensions (creating a genetic “line of least resistance” for the evolution of a population; Schluter 1996). This partly alleviates the apparent tension of robustness and evolvability. ii) But more interestingly, evolved correlations, uniquely, can reduce the dimensionality of phenotype space by reducing the *combinations* of traits that occur in phenotypes, without reducing phenotypic variability in any *individual* phenotypic characters. For example, the two phenotypes ++ and -- constitute only half of the four possible phenotypes composed of two bi-allelic characters, but each individual trait can still take either value + or - (this is the developmental analogue of altering the linkage disequilibrium of simultaneously segregating alleles without altering allele frequencies or fixing alleles at either locus). Limiting character combinations without limiting the variability of the individual characters alleviates the tension between robustness and evolvability in a more profound sense.

The logic functions IFF and XOR correspond to positive and negative correlation, respectively. These functions are significant in machine learning because the output is not a linearly decomposable function of the inputs (McClelland & Rumelhart 1986). It is necessary to learn correlations to represent these functions, and representing these functions is necessary to produce a memory that can store more than one pattern without simply blending them. Simpler types of canalisation, e.g. from a linear developmental process, preclude this possibility (as illustrated in Figs. 3.h & 4.d).

(j) Recall from partial stimuli showing developmental time

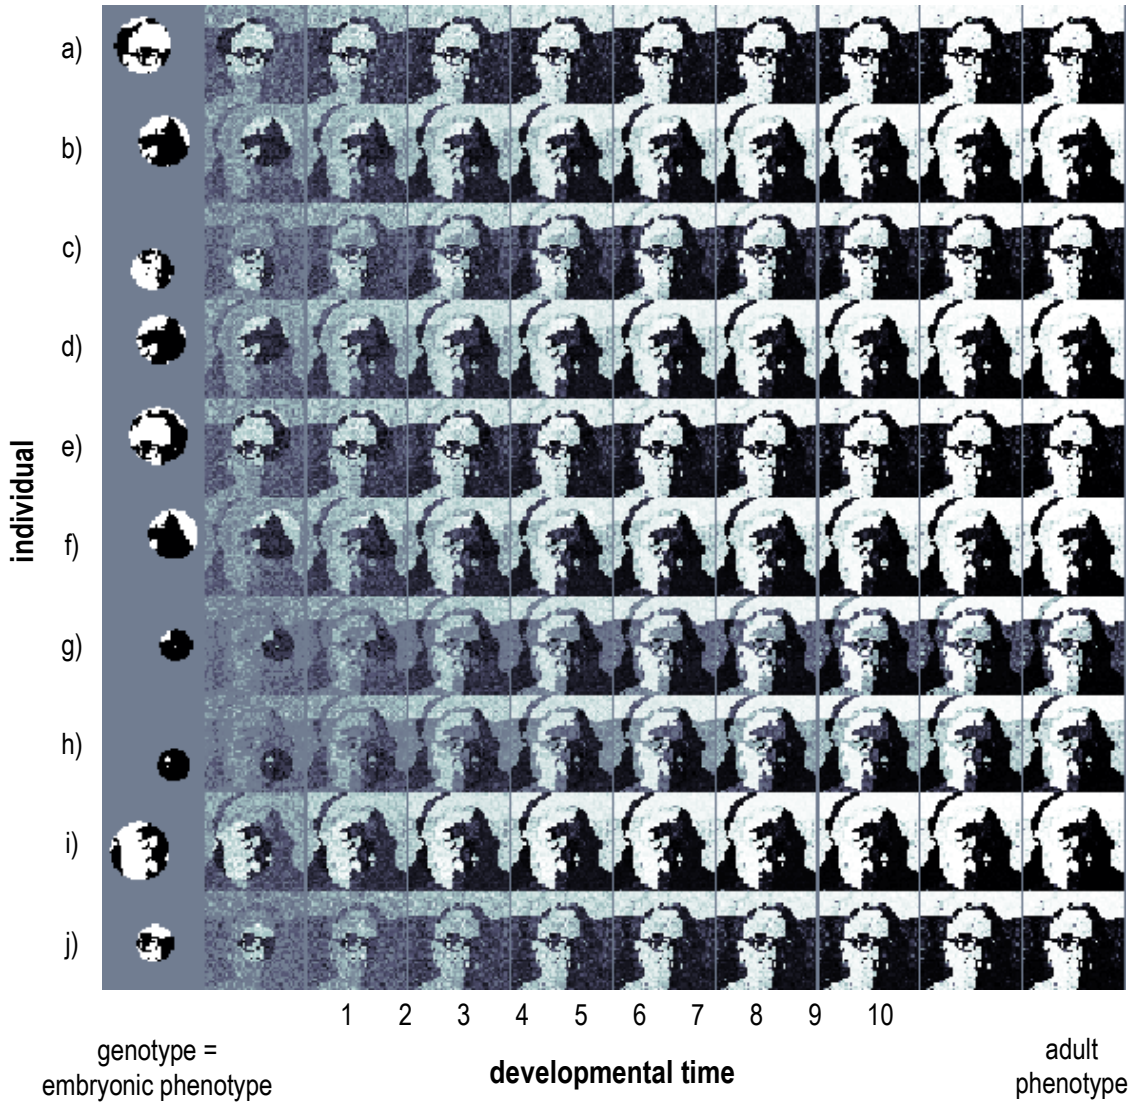

Fig. S5: An evolved gene regulation network exhibits a developmental memory of phenotypes that have been selected for in the past. The gene network was evolved in an environment where selection varies between two target phenotypes; an image of Charles Darwin and an image of Donald Hebb (the neuropsychologist from whom Hebbian learning takes its name). Each pixel in the image corresponds to the expression level of a gene. The evolved network of regulatory interactions introduces phenotypic correlations that create developmental attractors corresponding to these two phenotypes. We find that this is functionally equivalent to the manner in which a neural network learns an associative memory of a set of training patterns via Hebbian learning. This experiment shows that if the genotype, controlling the embryonic phenotype, partially resembles one of the targets, development 'recognises' which phenotype it belongs to and recreates the complete adult phenotype in a self-consistent manner. The development of ten example individuals from different genotypes (a-j) are shown over ten developmental time steps. In most cases development produces an adult phenotype that fully resembles only one of the phenotypes selected in the past. In cases where the genotype is almost equally matched to both phenotypes (e.g. g and h) the process of development breaks this symmetry, producing an adult phenotype that largely resembles only one, but some residual trace of the other phenotype is visible late into developmental time.

Literature Cited

See also main text.

Fontanari, J. F. 1990 Generalization in a Hopfield network. *J. Phys*. 51:2421–2430.

Hopfield, J. J., and D.W. Tank. 1986. Computing with neural circuits: a model. *Science* 233:625–633.

Jang, J.-S., M. W. Kim, and Y. Lee. 1992. A conceptual interpretation of spurious memories in the Hopfield-type neural network. *Neural Networks, 1992. IJCNN., International Joint Conference on* 1:21–26.

Mills, R. 2010. *How micro-evolution can guide macro-evolution: multi-scale search via evolved modular variation.* Ph.D. thesis, ECS, University of Southampton, Southampton, U.K.

Mills, R., and R. A. Watson. 2009. Symbiosis enables the evolution of rare complexes in structured environments. Pp. 110–117 in *Proceedings of 10th European Conference on Artificial Life*.

Mills, R., T. Jansen, and R. A. Watson. (under submission) *Transforming evolutionary search into higher-level evolutionary search by capturing problem structure*.