Published in final edited form as:

Evolution. 2014 April ; 68(4): 1124–1138. doi:10.1111/evo.12337.

The evolution of phenotypic correlations and ‘developmental memory’

Richard A. Watson¹, Günter P. Wagner², Mihaela Pavlicev³, Daniel M. Weinreich⁴, and Rob Mills⁵

Richard A. Watson: R.A.Watson@soton.ac.uk; Günter P. Wagner: gunter.wagner@yale.edu; Mihaela Pavlicev: mihaela.pavlicev@univie.ac.at; Daniel M. Weinreich: daniel_weinreich@brown.edu; Rob Mills: Rob.Mills@soton.ac.uk

¹Natural Systems Group, ECS/Institute for Life Sciences/Institute for Complex Systems Simulation, University of Southampton. SO17 1BJ. U.K

²Ecology and Evolutionary Biology, Yale University. CT 06477. USA

³Theoretical Biology, University of Vienna, Austria

⁴Ecology and Evolutionary Biology/Center for Computational Molecular Biology, Brown University, RI 02912. USA

⁵ECS, University of Southampton. SO17 1BJ. U.K

Abstract

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.

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).

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 multimodal. 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.

¹Other 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.

Hebb's rule simply states that the change in strength of a synaptic connection, w_{ij} , is proportional to the co-activation of the neurons it connects: $w_{ij} = r s_i s_j$, 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 phenotypic characters or traits, naturally represented by a vector, $P(t) = \langle p_1(t), p_2(t), \dots, p_N(t) \rangle$, $p \in \mathbb{R}$. The genotype of an individual has two parts, naturally represented by a vector of direct effects on traits, $G = \langle g_1, g_2, \dots, g_N \rangle$, 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 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,

$$P(t+1)=P(t)+\tau_1\sigma(B \times P(t))-\tau_2P(t), \quad \text{Eq. 1}$$

where $\tau_1=1$ is a rate constant controlling the magnitude of the interaction terms, $\tau_2=0.2$ 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 $P = G + \tau_1(B \times G) - \tau_2G$, 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:

$$p_i(t+1)=p_i(t)+\tau_1\sigma\left(\sum_{j=0}^n b_{ij}p_j(t)\right)-\tau_2p_i(t), \quad \text{Eq. 2}$$

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, $\bar{P}^*$, 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, $S = \langle s_1, s_2, s_3, \dots, s_N \rangle$, $s \in \{-1, 1\}$, 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 below).

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, $w_{ij} = r s_i s_j$ ($r > 0$) (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.

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., $s_1s_2 > 0$ in $S1$ and $s_1s_2 < 0$ 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.

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).

These phenotypes have $N=1,764$ genes (42×42 pixels) and accordingly, if the network were fully-connected, the number of interaction coefficients in B would be $N^2 = 3.1 \cdot 10^6$. 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 multilayer 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 nonlinear 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 the developmental memory demonstrated above 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). 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).

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., $b_{ij} = r s_i s_j$, ($r > 0$). In general, the selection coefficient of a mutation will be given by

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

where $\mathbf{P}^*$ 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 b_{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 b_{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 $\frac{\Delta p_i s_i}{1 + p_i s_i} = \frac{\tau_1 \Delta b_{ij} p_j s_i}{1 + p_i s_i}$. Given that p_i and s_i have the same sign, this mutation is favoured if and only if $\tau_1 b_{ij} p_j s_i > 0$. Given $\tau_1 > 0$, a change

b_{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, $b_{ij}/s_i p_j = r$, where r is some positive parameter. Thus all beneficial mutations satisfy the condition $b_{ij} = r s_i p_j$, ($r > 0$). (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, $b_{ij} = r s_i s_j$, ($r > 0$), 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 b_{ij} can then percolate into additional components of P^*), simulations show that the same effect obtains using the multi-step nonlinear 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. submitted-b), 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, submitted-a).

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. 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). By making a formal connection between the evolution of development and the principles of memory and learning we provide access to a theoretical framework for tackling these questions. In particular, the principles of generalisation, such as the tension between short-term performance and long-term 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 manner (Watson et al., submitted-b). Specifically, the equivalence of the selective pressures on ontogenetic interactions with associative learning mechanisms demonstrated here offers a specific mechanism by which 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.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

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.

References

- Ackley DH, Hinton GE, Sejnowski TJ. A Learning Algorithm for Boltzmann Machines. *Cognitive Science*. 1985; 9:147–169.
- Brigandt I. Typology now: homology and developmental constraints explain evolvability. *Biol Philos* (2007). 2007; 22:709–725.
- Cheverud JM, Ehrich TH, Vaughn TT, Koreishi SF, Linsey RB, Pletscher LS. Pleiotropic effects on mandibular morphology II: differential epistasis and genetic variation in morphological integration. *Journal of Experimental Zoology Part B: Molecular and Developmental Evolution*. 2004; 302(5): 424–435.
- Chevillon C, Bourget D, Rousset F, Pasteur N, Raymon M. Pleiotropy of adaptive changes in populations: comparisons among insecticide resistance genes in *Culex pipiens*. *Genet Res*. 1997:195–204. [PubMed: 9494436]
- Clune J, Mouret JB, Lipson H. The evolutionary origins of modularity. *Proceedings Royal Society B*. 2013; 280(1755)

- Coolen ACC. On the Relation Between Stability Parameters and Sizes of Domains of Attraction in Attractor Neural Networks. *Europhys Lett.* 1991; 16:73.
- Crombach A, Hogeweg P. Evolution of Evolvability in Gene Regulatory Networks. *PLoS Comput Biol.* 2008; 4(7)
- Davidson EH. Emerging properties of animal gene regulatory networks. *Nature.* 2010; 468(7326):911–920. [PubMed: 21164479]
- Davies AP, Watson RA, Mills R, Buckley CL, Noble J. 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.* 2010; 17 (3):167–181. [PubMed: 21554113]
- Delph LF, Steven JC, Anderson IA, Herlihy CR, Brodie ED III. Elimination of a Genetic Correlation Between the Sexes Via Artificial Correlational Selection. *Evolution.* 2011; 65(10):2872–2880. [PubMed: 21967428]
- Draghi J, Wagner GP. Evolution of evolvability in a developmental model. *Evolution.* 2008; 62:301–315. [PubMed: 18031304]
- Draghi J, Wagner GP. The evolutionary dynamics of evolvability in a gene network model. *Journal of Evolutionary Biology.* 2009; 22(3):599–611. [PubMed: 19170816]
- Draghi JA, Parsons TL, Wagner GP, Plotkin JB. Mutational robustness can facilitate adaptation. *Nature.* 2010; 463(7279):353–355. [PubMed: 20090752]
- Fernando C, Liekens AML, Bingle LEH, Beck C, Lenser T, Stekel DJ, Rowe JE. Molecular circuits for associative learning in single-celled organisms. *Journal of the Royal Society Interface.* 2008; 6:463–469.
- Fontanari JF. Generalization in a Hopfield network. *Journal de Physique.* 1990; 51:2421–2430.
- Gillespie JH. Molecular evolution over the mutational landscape. *Evolution.* 1984; 38:1116–1129.
- Guillaume F, Otto SP. Gene Functional Trade-offs and the Evolution of Pleiotropy. *Genetics.* 2012
- Gross, T.; Sayama, H. *Adaptive Networks. Theory, Models and Applications.* Berlin: Springer-Verlag; 2009. Series: Understanding Complex Systems
- Harnad, S. To Cognize is to Categorize: Cognition is Categorization. In: Lefebvre, C.; Cohen, H., editors. *Handbook of Categorization in Cognitive Science.* New York: Elsevier Press; 2005.
- Hebb, DO. *The organization of behaviour.* New York: Wiley; 1949.
- Hendrikse JL, Parsons TE, Hallgrímsson B. Evolvability as the proper focus of evolutionary developmental biology. *Evol Dev.* 2007; 9(4):393–401. [PubMed: 17651363]
- Hopfield JJ. Neural networks and physical systems with emergent collective computational abilities. *PNAS USA.* 1982; 79(8):2554–2558. [PubMed: 6953413]
- Hopfield JJ, Tank DW. Computing with neural circuits: A model. *Science.* 1986; 233:625–633. [PubMed: 3755256]
- Huang S, Eichler G, Bar-Yam Y, Ingber DE. Cell fates as high-dimensional attractor states of a complex gene regulatory network. *Physical review letters.* 2005; 94(12):128701. [PubMed: 15903968]
- Izquierdo EJ, Fernando C. The Evolution of Evolvability in Gene Transcription Networks. *Proceedings of Alife XI.* 2008
- Jang, J-S.; Kim, MW.; Lee, Y. A Conceptual Interpretation of Spurious Memories in the Hopfield-type Neural Network. *Neural Networks, 1992. IJCNN., International Joint Conference on; 1992.* p. 21–26.
- Jones AG, Arnold SJ, Bürger R. The mutation matrix and the evolution of evolvability. *Evolution.* 2007; 61:727–745. [PubMed: 17439608]
- Kashtan N, Alon U. Spontaneous evolution of modularity and network motifs. *PNAS USA.* 2005; 102(39):13773–13778. [PubMed: 16174729]
- Kashtan N, Mayo A, Kalisky T, Alon U. An Analytically Solvable Model for Rapid Evolution of Modular Structure. *PLoS Comput Biol.* 2009; 5(4)
- Kashtan N, Noor E, Alon U. Varying Environments Can Speed up Evolution. *Proceedings of the National Academy of Sciences of the United States of America.* 2007 Aug 21; 104(34):13711–13716. [PubMed: 17698964]
- Kauffman, SA. *The Origins of Order.* Oxford University Press; New York: 1993.

- Kim HS, Huh J, Fay JC. Dissecting the pleiotropic consequences of a quantitative trait nucleotide. *FEMS yeast research*. 2009; 9(5):713–722. [PubMed: 19456872]
- Kirchner M, Gerhart J. Evolvability. *Proc Natl Acad Sci USA*. 1998; 95:8420–8427. [PubMed: 9671692]
- Laland KN, Sterelny K, Odling-Smee J, Hoppitt W, Uller T. Cause and effect in biology revisited: is Mayr's proximate-ultimate dichotomy still useful? *Science*. 2011; 334(6062):1512–1516. [PubMed: 22174243]
- Lande R, Arnold SJ. The measurement of selection on correlated characters. *Evolution*. 1983; 37:1210–1226.
- Leamy LJ, Pomp D, Lightfoot JT. Genetic variation in the pleiotropic association between physical activity and body weight in mice. *Genet Sel Evol*. 2009; 41:41. [PubMed: 19775457]
- Leclerc RD. Survival of the sparsest: robust gene networks are parsimonious. *Molecular systems biology*. 2008; 4(1)
- Lenski RE. Experimental studies of pleiotropy and epistasis in *Escherichia coli*. I. Variation in competitive fitness among mutants resistant to virus T4. *Evolution*. 1988a:425–432.
- Lenski RE. Experimental studies of pleiotropy and epistasis in *Escherichia coli*. II. Compensation for maladaptive effects associated with resistance to virus T4. *Evolution*. 1988b:433–440.
- Lewis, M. MSc Dissertation. ECS, University of Southampton; 2009. An Investigation Into The Evolution Of Relationships Between Species In An Ecosystem.
- Lipson H. Principles of modularity, regularity, and hierarchy for scalable systems. *Journal of Biological Physics and Chemistry*. 2007; 7(4):125–128.
- Lipson H, Pollack JB, Suh NP. On the Origin of Modular Variation. *Evolution*. 2002 Aug; 56(8): 1549–1556. [PubMed: 12353747]
- McGregor S, Vasas V, Husbands P, Fernando C. Evolution of Associative Learning in Chemical Networks. *PLoS Computation Biology*. 2012; 8:11, e1002739.
- Mills, R. PhD thesis. ECS, University of Southampton; 2010. How Micro-Evolution Can Guide Macro-Evolution: Multi-Scale Search via Evolved Modular Variation.
- Mills R, Jansen T, Watson RA. Transforming Local Search Into Higher-Order Search by Capturing Problem Structure. submitted.
- Mills, R.; Watson, RA. Symbiosis Enables the Evolution of Rare Complexes in Structured Environments. *Proceedings of 10th European Conference on Artificial Life (ECAL 2009)*; 2009. p. 110-117.
- O'Reilly, RC.; Munakata, Y. *Computational Explorations in Cognitive Neuro-science: Understanding the Mind by Simulating the Brain*. MIT Press; Cambridge, MA: 2000.
- Parter M, Kashtan N, Alon U. Facilitated Variation: How Evolution Learns from Past Environments To Generalize to New Environments. *PLoS Comput Biol*. 2008; 4(11):e1000206. [PubMed: 18989390]
- Pavlicev M, Wagner GP. A model of developmental evolution: selection, pleiotropy and compensation. *Trends in Ecology & Evolution*. 2012; 27(6):316–322. [PubMed: 22385978]
- Pavlicev M, KenneyHunt JP, Norgard EA, Roseman CC, Wolf JB, Cheverud JM. Genetic variation in pleiotropy: differential epistasis as a source of variation in the allometric relationship between long bone lengths and body weight. *Evolution*. 2008; 62 (1):199–213. [PubMed: 18005158]
- Pavlicev M, Cheverud JM, Wagner GP. Evolution of adaptive phenotypic variation patterns by direct selection for evolvability . *Proceedings of the Royal Society B: Biological Sciences*. 2011; 278(1713):1903–1912.
- Raff RA. Evo-devo: the evolution of a new discipline. *Nature Rev Genet*. 2000; 1:74–79. [PubMed: 11262880]
- Riedl, R. *Order in Living Organisms*. Wiley; 1978.
- Rumelhart, DE.; James, L. McClelland and the PDP Research Group. *Parallel Distributed Processing*. Vol. I and II. MIT Press; 1986.
- Sansom, R. *Ingenious Genes: How Gene Regulation Networks Evolve to Control Development*. MIT Press; 2011.
- Schluter D. Adaptive radiation along genetic lines of least resistance. *Evolution*. 1996; 50:1766–1774.

- Sniegowski PD, Murphy HA. Evolvability. *Current Biology*. 2006; 16:R831–R834. [PubMed: 17027474]
- Stewart J. The evolution of genetic cognition. *Journal of Social and Evolutionary Systems*. 1997; 20(1):53–73.
- Toussaint M, von Seelen W. Complex adaptation and system structure. *BioSystems*. 2007; 90:769–782. [PubMed: 17512656]
- Vohradsky J. Neural model of the genetic network. *The Journal of Biological Chemistry*. 2001a; 276:36168–36173. [PubMed: 11395518]
- Vohradsky J. Neural network model of gene expression. *The FASEB journal*. 2001b; 15(3):846–854.
- Waddington CH, Robertson E. Selection for developmental canalisation. *Genet Res*. 1966; 7:303–312. [PubMed: 5940870]
- Wagner GP, Laubichler MD. Rupert Riedl and the re-synthesis of evolutionary and developmental biology: body plans and evolvability. *J Exp Zool B: Mol Dev Evol*. 2004; 302B:92–102. [PubMed: 14760655]
- Wagner GP. Multivariate mutation-selection balance with constrained pleiotropic effects. *Genetics*. 1989; 122:223–234. [PubMed: 2731730]
- Wagner GP, Altenberg L. Complex adaptations and the evolution of evolvability. *Evolution*. 1996; 50:967–976.
- Wagner GP, Pavlicev M, Cheverud JM. The road to modularity. *Nature Reviews Genetics*. 2007; 8:921–931.
- Watson, RA.; Palmius, N.; Mills, R.; Powers, ST.; Penn, AS. Can Selfish Symbioses Effect Higher-level Selection?. *Proceedings of 10th European Conference on Artificial Life (ECAL 2009)*; 2009. p. 27-36.
- Watson RA, Buckley CL, Mills R, Davies AP. Associative memory in gene regulation networks. *ALIFE XII*. 2010a:194–202.
- Watson, RA.; Mills, R.; Buckley, CL.; Penn, A.; Davies, A.; Noble, J.; Bullock, S. Adaptation without natural selection. *Artificial Life XII: Proceedings of the Twelfth International Conference on the Synthesis and Simulation of Living Systems*; MIT Press; 2010. p. 80-81.
- Watson RA, Szathmari E, Power D, Mills R, Powers ST, Czapp B. The Evolution of Ecological Relationships and Connectionist Principles of Adaptation in Ecosystems. in prep. Manuscript in preparation.
- Watson, RA.; Buckley, CL.; Mills, R. Technical Report. ECS, University of Southampton; 2009. The Effect of Hebbian Learning on Optimisation in Hopfield Networks.
- Watson RA, Buckley CL, Mills R. Optimisation in ‘Self-modelling’ Complex Adaptive Systems. *Complexity*. 2010; 16(5):17–26.
- Watson RA, Mills R, Buckley CL. Global Adaptation in Networks of Selfish Components: Emergent Associative Memory at the System Scale. *Artificial Life*. 2010; 17(3):147–166. [PubMed: 21554114]
- Watson RA, Mills M, Buckley CL. Transformations in the Scale of Behaviour and the Global Optimisation of Constraints in Adaptive Networks. *Adaptive Behavior*. 2011; 19(4):227–249.
- Watson RA, Weinreich DM, Wakeley J. Genome Structure and the Benefit of Sex. *Evolution*. 2011; 65(2):523–536. [PubMed: 21029076]
- Watson RA, Jackson A, Palmius N, Mills RM, Powers ST. The Evolution of Symbiotic Partnerships, the Adaptive Capabilities of Ecosystems and the Algorithmic Principles underlying the Major Evolutionary Transitions. submitted-a.
- Watson RA, Mills R, Buckley CL, Kouvaris K. Evo-Connectionism I: The Associative Learning Principles of Evolutionary Systems. submitted-b.
- Wessels L, van Someren E, Reinders M. A comparison of genetic network models. In *Pacific Symposium on Biocomputing*. 2001; 6:508–519.
- West-Eberhard, M-J. *Developmental plasticity and evolution*. Oxford University Press; New York: 2003.
- Wimsatt WC. Generative entrenchment and the developmental systems approach to evolutionary processes. *Cycles of contingency: Developmental systems and evolution*. 2001:219–37.

- Young NM, Hallgrímsson B. Serial homology and the evolution of mammalian limb covariation structure. *Evolution*. 2005; 59(12):2691–2704. [PubMed: 16526515]
- Young NM, Wagner GP, Hallgrímsson B. Development and the evolvability of human limbs. *Proceedings of the National Academy of Sciences*. 2010; 107(8):3400–3405.

Fig. 1.

Experiment 1. A system of 8 genes, selecting for a single (arbitrary) phenotypic pattern ($SI = ++--+-+$). **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 $s_i s_j$. **c)** Interaction matrix derived from Hebb's rule (i.e., $w_{ij} = r s_i s_j$) 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.

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.

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).

Fig. 4.

Experiment 4. Modularly-varying environment. a) Target phenotypes, $N=100$ genes (10×10 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.