doi:10.1111/evo.12337

THE EVOLUTION OF PHENOTYPIC CORRELATIONS AND "DEVELOPMENTAL MEMORY"

Richard A. Watson,[1,2] Günter P. Wagner,[3] Mihaela Pavlicev,[4] Daniel M. Weinreich,[5] and Rob Mills[6]

[1]*Natural Systems Group, ECS/Institute for Life Sciences/Institute for Complex Systems Simulation, University of Southampton, Southampton SO17 1BJ, United Kingdom*

[2]*E-mail: R.A.Watson@soton.ac.uk*

[3]*Ecology and Evolutionary Biology, Yale University, Connecticut 06477*

[4]*Theoretical Biology, University of Vienna, Vienna, Austria*

[5]*Ecology and Evolutionary Biology/Center for Computational Molecular Biology, Brown University, Providence, Rhode Island*

[6]*ECS, University of Southampton, Southampton SO17 1BJ, United Kingdom*

Received May 23, 2013
Accepted November 22, 2013

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 nonlinear 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 "generalize" (by exploiting evolved developmental modules) to produce new combinations of phenotypic features. We show that these surprising behaviors 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.

KEY WORDS: **Adaptation, associative learning, evolvability, evo-devo.**

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 covary, creating patterns of phenotypic variation that are thereby partially nonrandom. Because developmental processes are themselves a product of evolution, such biases and constraints can, in principle, be shaped by past selection (Riedl 1978; Raff 2000; Brigandt 2007; Hendrikse et al. 2007; Wagner et al. 2007; Crombach and Hogeweg 2008; Izquierdo and Fernando 2008; Draghi et al. 2010; Pavlicev and Wagner 2012; Clune et al. 2013).

We seek general organizational principles to understand how past selective environments can alter phenotypic correlations and

EVOLUTION
INTERNATIONAL JOURNAL OF ORGANIC EVOLUTION

hence shape the distribution of phenotypic variants produced by development in adaptive ways (Toussaint and von Seelen 2007; Wagner et al. 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 macroscale morphological growth. Heritable genetic variation affecting phenotypic correlations has been shown in quantitative data from many organisms (Lenski 1988a, b; Chevillon et al. 1997; Cheverud et al. 2004; Pavlicev et al. 2008; Kim et al. 2009; Leamy et al. 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 hind–limb correlations in mammals (Young et al. 2005) and in primates in particular (Young et al. 2010). Characterizing 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 and Otto 2012). To begin to explain these patterns, Pavlicev et al. (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, that is, cause subsets of traits to vary together but independently of other subsets (Wagner et al. 2007). Early work on this idea (Lipson et al. 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 favors 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 utilize 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 among the multiple features of sepals and petals, but the resultant distribution also includes phenotypes all along the (multidimensional) line in between sepals and petals.

In principle, a nonlinear model of a developmental process or a more complex G-P mapping could represent and/or produce multimodal phenotype distributions. For example, Kashtan et al. (2007, also Kashtan and 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 et al. (2008) develop analogies with memory, that is, 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 "generalize 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 nonlinear developmental processes?

Here we build on this prior work with the aim of identifying organizational principles to predict how past selection shapes the properties of nonlinear 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 behaviors 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 nonlinear model of ontogenetic interactions. Some nonlinearity in the mapping is important (as we will show) but this can be of a simple, natural

form; for example, 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 nonlinear interactions transform a set of embryonic phenotypic characters into their adult form over multiple developmental time steps.

We find that this nonlinear 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 multimodal 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 generalize past phenotypic patterns, for example, by evolving developmental modules and producing new phenotypes that are built from novel combinations of those modules.

In addition to illustrating these memory behaviors 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 nonlinear 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 behaviors. 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 behaviors 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 behaviors.

Second, Pavlicev et al. (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 favors an increase in correlation (or decrease in anticorrelation) of those traits. Conversely, if the traits are selected contrariwise, that is one is under positive directional selection and the other negative, then selection favors anticorrelation (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 et al. 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 et al. 1985; Rumelhart et al. 1986; O'Reilly and Munakata 2000).[1] *Hebb's rule* simply states that the change in strength of a synaptic connection, Δw_{ij}, is proportional to the coactivation of the neurons it connects: $\Delta w_{ij} = r s_i s_j$, 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 et al.'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 behaviors we observe in an evolved network of recurrent nonlinear interactions. That is, a gene network can evolve regulatory interactions that "internalize" a model of past selective environments in just the same way that a learning neural network can store, recall, recognize, and generalize a set of training patterns. Recognizing this equivalence between associative learning and evolution of phenotypic correlations thus provides access to an established theoretical framework that we can use to characterize organizational principles describing how regulatory interactions evolve; including the affordances and limitations of developmental memory, the minimal conditions for such behaviors and their potential impact on evolvability. It is a main aim of this article to explain the dynamical and functional equivalence of these evolved developmental behaviors 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 behaviors.

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

[1]Other work has investigated the potential to implement associative learning mechanisms in various nonneural 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.

represented by a vector, $P(t) = <p_1(t), p_2(t), \ldots, p_N(t)>$, $p \in \mathbb{R}$. The genotype of an individual has two parts, naturally represented by a vector of direct effects on traits, $G = <g_1, g_2, \ldots, g_N>$, and the elements b_{ij} of an interaction matrix, B (Lande and Arnold 1983; Wagner 1989; Lipson et al. 2002; Jones et al. 2007; Kashtan et al. 2009). Although previous work utilized a linear model of phenotypic correlations where B^T was a genetic covariance matrix among the adult characters, here we analyze a nonlinear model of a developmental process, where b_{ij} represents the interaction coefficient between trait i and trait j within a dynamical ontogenetic process modeled as follows.

Let the initial embryonic phenotype of an organism at developmental time $t = 0$, that is $P(0)$, 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 nonlinear transform determined by the matrix B, that is the weighted influence of each character on each other character, and a decay term. Specifically, in linear algebra form,

$$P(t+1) = P(t) + \tau_1 \sigma(B \times P(t)) - \tau_2 P(t), \quad (1)$$

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 nonlinearly 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 et al. (2002) and Kashtan et al. (2009). In the absence of interactions, that is if all off-diagonal elements of B are 0, equation (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 nonlinear transformation has analogues in many natural systems and biological processes. An example of ontogenetic process that can be modeled 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 et al. 2001), which develops the "embryonic" activity levels of each gene into an "adult" pattern of activity. As per equation (1), the new activity level of a gene, $p_i \in P$, in a single time step is given by

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

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 nonlinear 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, that is $P^* = P(t = T)$. 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, that is, $P^* = develop(G,B,T)$.

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 and 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 emphasize this point here we show that strong selection weak mutation (SSWM) assumptions (Gillespie 1984), that is 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. 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 et al. (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, $S = <s_1, s_2, s_3, \ldots, s_N>$, $s \in \{-1,1\}$, 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, that is $w(P^*) = 1 + P^*S$. The elements of S thus determine the direction of selection on each of the phenotypic characters (Pavlicev et al. 2011).

Lipson et al. (2002), Kashtan and Alon (2005), Kashtan et al. (2007, 2009), Draghi and Wagner (2008, 2009) and Pavlicev et al. (2011) all investigate evolution in varying selective environments. This makes intuitive sense because stabilizing selection in a single environment offers no advantage to genotypes that admit variability, structured or otherwise, and will simply tend to canalize phenotypes (Waddington and Robertson 1966), reducing phenotypic variability in all dimensions. In contrast, directional selection can favor 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 utilizing 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 because an individual only needs to exhibit one phenotype to maximize 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 predisposition 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 initialized to 0. 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, that is, $|g| \leq 1$. 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 Figure 1(A) and (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 et al. 2011), that is direction of selection

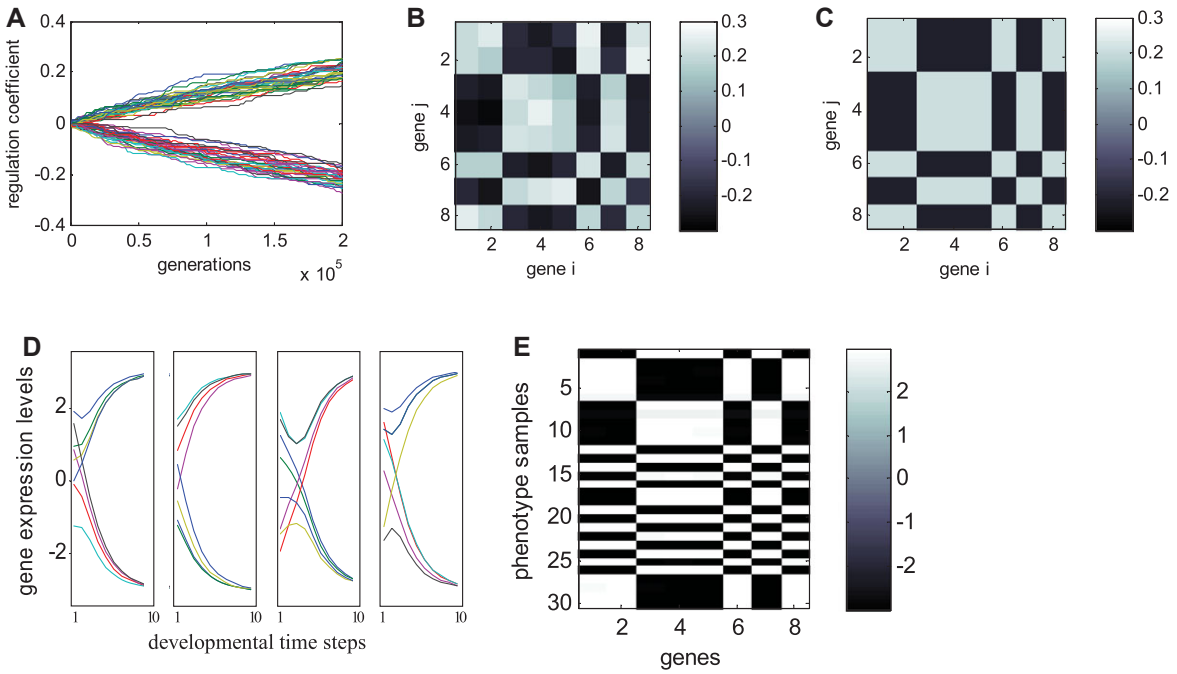

Figure 1. **Experiment 1. A system of eight genes, selecting for a single (arbitrary) phenotypic pattern (*S1* = ++−−−+−+). (A) Regulatory interaction coefficients evolve into positive and negative classes. (B) The matrix of evolved regulatory interactions in the gene regulation network (at generation 2×10^5). 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) Thirty independent adult phenotypes developed from random G. Development produces either the target phenotype or its complement.**

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 favors 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 utilized, 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, although 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 emphasize that this is also in agreement with Hebb's rule, $\Delta w_{ij} = r s_i s_j$ ($r > 0$) (see *Analysis*). Figure 1(B) shows the evolved interactions, whereas Figure 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 et al. 2011). Figure 1(D) and (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 canalization.

EXPERIMENTS 2 AND 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. Figure 2(A) shows that evolved interactions fall into three classes. Some evolve at a constant positive rate; these arise from

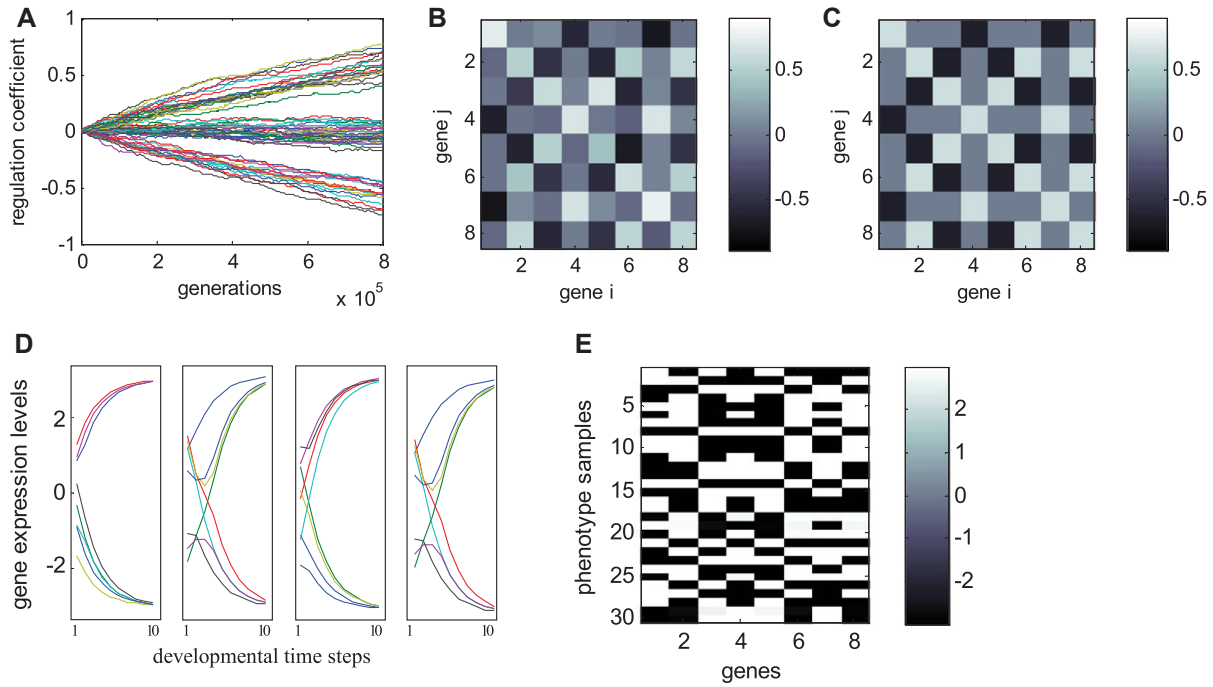

Figure 2. Experiment 2. Evolved interaction coefficients for changing environment (two target patterns *S1* = ++−−−+−+, *S2* = +−+−+−−−). Target patterns alternate each 2000 generations. (A) Interaction coefficients over evolutionary time fall into three classes. (B) Evolved interaction matrix after 8×10^5 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.

pairs of traits that are positively correlated in *both* patterns (e.g., genes 2 and 6 are ++ in *S1* and −− in *S2*), similarly 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). Figure 2(B) and (C) show 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 (Fig. 2B, C), and that the effect of these evolved correlations is that development reproduces either of the two target patterns (Fig. 2E). 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 recognizable 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. 3A, B).

These phenotypes have $N = 1764$ genes (42×42 pixels) and accordingly, if the network were fully-connected, the number of interaction coefficients in B would be $N^2 = 3.1\times10^6$. 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., Leclerc 2008; Davidson 2010). 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

Figure 3. Experiment 3. (A) Target phenotype *S1* = image of Charles Darwin; each pixel indicates whether a given gene should be upregulated (white) or downregulated (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 0); 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).

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. 3C, D). Notice that the output of the network is extremely robust to mutation in G (Fig. 3C), far beyond what is easily recognizable 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. 3D), 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). Figure 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 Figure 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 NONLINEAR DEVELOPMENTAL INTERACTIONS

For the most part, this nonlinear developmental mapping produces adult phenotypes *S1* or *S2* without "mangling" the two patterns—that is, 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. 3F and 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. 3G).

Figure 3(H) shows the distribution of (left) target phenotypes, (center) the evolved adult phenotypes, and (right) phenotypes produced by the linear G-P map. The phenotypic distribution from the nonlinear 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. 3F) 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 canalized 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—that is, to the extent that such a distribution remembers the most recently selected phenotype all memories of previously 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 because only a nonlinear developmental process is capable of holding more than one memory.

The significance of nonlinearity in the input–output transform function and its implications with respect to these behaviors is well understood in the context of neural networks (both for feed-forward multilayer networks and recurrent networks; Hopfield 1982; Rumelhart et al. 1986). 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 nonlinear (this is related to the fact that a single-layer nonlinear perceptron cannot represent XOR; Rumelhart et al. 1986). In contrast, with a nonlinear 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 behavior is referred to as "recognition" or "categorical perception" (Harnad 2005)—that is, identifying which discrete class an input pattern belongs to (Fig. 3F).

EXPERIMENT 4: STRUCTURED VARIATION: GENERALIZATION 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 (Fig. 4) assesses the propensity of a developmental memory to generalize over a collection of phenotypic targets and produce novel phenotypes from the same class. The targets are

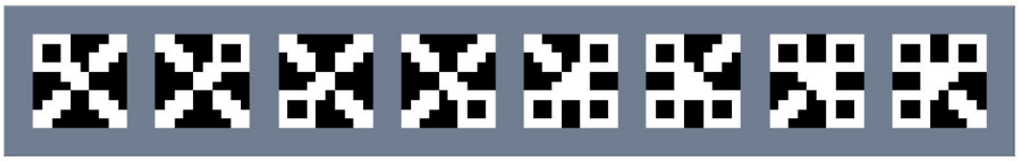

A Eight target phenotypes.

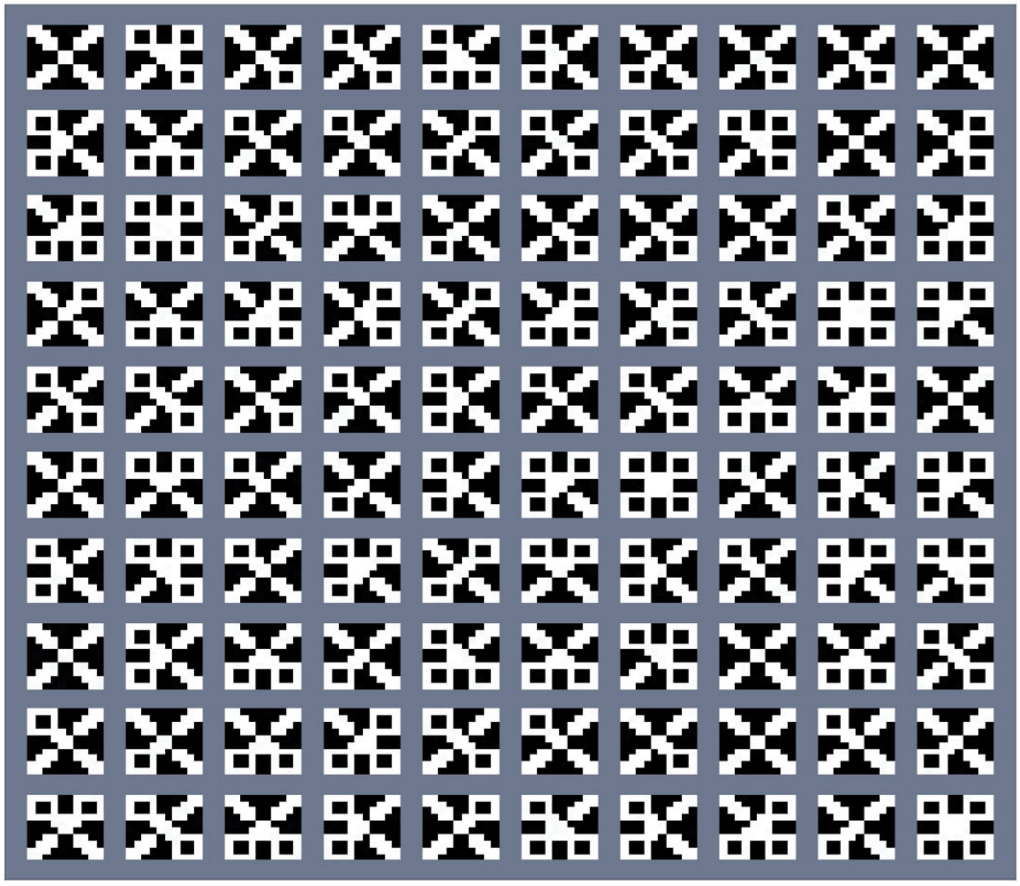

B Bestiary of 100 adult phenotypes produced by evolved developmental network from random *G*.

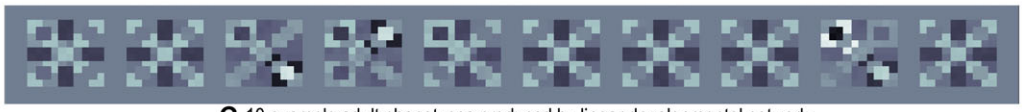

C 10 example adult phenotypes produced by linear developmental network.

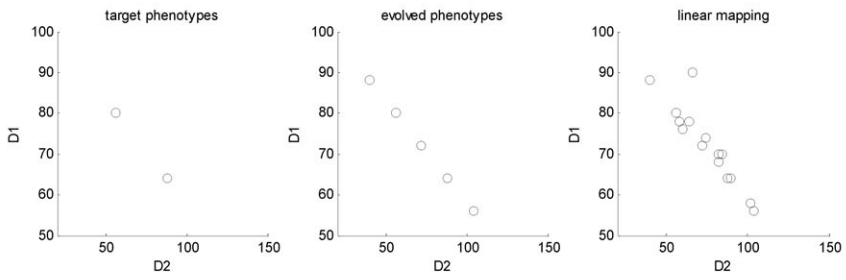

D Phenotype distributions of targets and adults with non-linear and linear development.

Figure 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 generalizations. (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.

drawn from a "family" of related patterns (Parter et al. 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 Figure 4(A). Each pattern has four subgroups of characters (the quartiles of each image) that vary in a simple modular fashion; that is, characters within the same subgroup are strongly correlated whereas correlations between these groups are weak or absent (Lipson et al. 2002; Lipson 2007; Wagner et al. 2007; Kashtan et al. 2009; Watson et al. 2011b; Clune et al. 2013). In this example, each subpattern (a "loop" or a "stalk" in various orientations) appears in at least one of the target phenotypes. Again, the particular subpatterns 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 et al. 2008). Here B is fully connected (not sparse).

Figure 4(B) examines the distribution of adult phenotypes produced from the evolved developmental interactions by randomizing G, as before. We see that development produces each of the eight training patterns, but it also produces simple generalizations of these phenotypes from new combinations of modules, for example, phenotypes with four loops. Figure 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 Figure 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 generalizes 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 generalize in both an interpolative and extrapolative manner from phenotypes that have been selected in the past. Supplementary text (f) discusses how this type of generalization is (necessarily) equivalent to a "failure" to restrict phenotypes to a set of training patterns accurately. Again the linear mapping (Fig. 4D, 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 internalizing 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 et al. 2008). In a different example, the modularity in the evolved interaction matrix that enables this generalization 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, that is, $\Delta b_{ij} = r s_i s_j$ $(r > 0)$. In general, the selection coefficient of a mutation will be given by

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

where ΔP^* is the vector of changes conferred on the adult phenotype. For a recurrent nonlinear 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 eq. 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

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

Thus, the selective coefficient of the mutation is $\frac{\Delta p_i s_i}{1 + p_i s_i} = \frac{\tau_1 b_{ij} p_j s_i}{1 + p_i s_i}$. Given that p_i and s_i have the same sign, this mutation is favored 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 minimization, 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 modeled 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 multistep 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 behaviors 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, that is, 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; that is, increase the number of initial conditions (here, embryonic phenotypes) that lead to that configuration (Watson et al. 2010c, d; 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 (Davies et al. 2010; Watson et al. 2010d, 2011a, 2014), including the tendency of natural selection on nontrophic ecological relationships to reinforce correlations in species densities and thereby produce "ecosystem memory" (Lewis 2009; Watson et al. 2011c, 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 multimodal phenotypic distribution, robustness goes hand-in-hand with sensitivity—that is, 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 Figure 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 article, 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, categorize patterns from partial or corrupted stimuli, and produce generalized patterns from a set of structurally similar training patterns has been well studied (e.g., Hopfield 1982; Rumelhart et al. 1986; O'Reilly and Munakata 2000). The insight that the selective pressures on developmental correlations are equivalent to associative learning thus provides the opportunity to utilize well-established theoretical and conceptual frameworks from associative learning theory to identify organizational principles involved in the evolution of development; for example, to understand the minimal conditions for a developmental memory capable of the behaviors 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 behaviors is more than merely a metaphor, and helps us make sense of how biological networks evolve adaptive complexity (Stewart 1997; Sansom 2011).

This article has focused 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 et al. 2007; Parter et al. 2008). Exactly how the evolved properties of development affect such evolvability (Wagner and Altenberg 1996; Kirchner and Gerhart 1998; Wagner and Laubichler 2004; Hendrikse et al. 2007; Wagner et al. 2007; Draghi and Wagner 2008, 2009; Izquierdo and Fernando 2008; Laland et al. 2011; Pavlicev and Wagner 2012) will be analyzed 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 generalization, such as the tension between short-term performance gains and long-term performance losses due to overfitting, may be useful for understanding the conditions for and limitations of evolvability (Supplementary text (f) and (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.

ACKNOWLEDGMENTS

The authors thank A. Davies, C. Buckley, R. Doursat, N. Barton, L. Kounios, T. Ezard, and J. Clune for discussion, and K. Kouvaris for technical contributions. This work in part supported by National Science Foundation Emerging Frontiers Award 1038657 to DMW.

LITERATURE CITED

Ackley, D. H., G. E. Hinton, and T. J. Sejnowski. 1985. A learning algorithm for Boltzmann machines, Cogn. Sci. 9:147–169.

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

Cheverud, J. M., T. H. Ehrich, T. T. Vaughn, S. F. Koreishi, R. B. Linsey, and L. S. Pletscher. 2004. Pleiotropic effects on mandibular morphology II: differential epistasis and genetic variation in morphological integration. J. Exp. Zool. Part B: Mol. Dev. Evol. 302:424–435.

Chevillon, C., D. Bourget, F. Rousset, N. Pasteur, and M. Raymon. 1997. Pleiotropy of adaptive changes in populations: comparisons among insecticide resistance genes in *Culex pipiens*. Genet. Res. 70:195–204.

Clune, J., J. B. Mouret, and H. Lipson. 2013. The evolutionary origins of modularity. Proc. R. Soc. B 280:20122863.

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., and P. Hogeweg. 2008. Evolution of evolvability in gene regulatory networks. PLoS Comput. Biol. 4:e1000112.

Davidson, E. H. 2010. Emerging properties of animal gene regulatory networks. Nature 468:911–920.

Davies, A. P., R. A. Watson, R. Mills, C. L. Buckley, and J. Noble. (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. Artif. Life 17:167–181.

Delph, L. F., J. C. Steven, I. A. Anderson, C. R. Herlihy, and E. D. Brodie, III 2011. Elimination of a genetic correlation between the sexes via artificial correlational selection. Evolution 65:2872–2880.

Draghi, J., and G. P. Wagner. 2008. Evolution of evolvability in a developmental model. Evolution 62:301–315.

———. 2009. The evolutionary dynamics of evolvability in a gene network model. J. Evol. Biol. 22:599–611.

Draghi, J. A., T. L. Parsons, G. P. Wagner, and J. B. Plotkin. 2010. Mutational robustness can facilitate adaptation. Nature 463:353–355.

Fernando, C., A. M. L. Liekens, L. E. H. Bingle, C. Beck, T. Lenser, D. J. Stekel, and J. E. Rowe. 2008 Molecular circuits for associative learning in single-celled organisms. J. R. Soc. Interface 6:463–469.

Gillespie, J. H. 1984. Molecular evolution over the mutational landscape. Evolution 38:1116–1129.

Guillaume, F., and S. P. Otto. 2012. Gene functional trade-offs and the evolution of Pleiotropy. Genetics 192:1389–1409.

Harnad, S. 2005. To cognize is to categorize: cognition is categorization. Pp. 19–43 *in* C. Lefebvre and H. Cohen. Handbook of categorization in cognitive science. Elsevier Press, New York.

Hebb, D. O. 1949. The organization of behaviour. Wiley, New York.

Hendrikse, J. L., T. E. Parsons, and B. Hallgrímsson. 2007. Evolvability as the proper focus of evolutionary developmental biology. Evol. Dev. 9:393–401.

Hopfield, J. J. 1982. Neural networks and physical systems with emergent collective computational abilities. Proc. Natl. Amer. Sci. USA 79:2554–2558.

Huang, S., G. Eichler, Y. Bar-Yam, and D. E. Ingber. 2005. Cell fates as high-dimensional attractor states of a complex gene regulatory network. Phys. Rev. Lett. 94:128701.

Izquierdo, E. J., and C. Fernando. 2008. The evolution of evolvability in gene transcription networks. Pp. 265–273 *in* S. Bullock, J. Noble, R. Watson, and M. A. Bedau, eds. International conference on the simulation and synthesis of living system. MIT Press, Cambridge, MA.

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

Kashtan, N., and U. Alon 2005. Spontaneous evolution of modularity and network motifs. Proc. Natl. Amer. Sci. USA 102:13773–13778.

Kashtan, N., E. Noor, and U. Alon. 2007. Varying environments can speed up evolution. Proc. Natl. Acad. Sci. USA 104:13711–13716.

Kashtan, N., A. Mayo, T. Kalisky, and U. Alon. 2009. An analytically solvable model for rapid evolution of modular structure. PLoS Comput. Biol. 5:e1000355.

Kauffman, S. A. 1993. The origins of order. Oxford Univ. Press, New York.

Kim, H. S., J. Huh, and J. C. Fay. 2009. Dissecting the pleiotropic consequences of a quantitative trait nucleotide. FEMS Yeast Res. 9: 713–722.

Kirchner, M., and J. Gerhart. 1998. Evolvability. Proc. Natl. Acad. Sci. USA 95:8420–8427.

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

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

Leamy, L. J., D. Pomp, and J. T. Lightfoot. (2009). Genetic variation in the pleiotropic association between physical activity and body weight in mice. Genet. Sel. Evol. 41:41. Available from http://www.gsejournal.org/about#citingarticles.

Leclerc, R. D. 2008. Survival of the sparsest: robust gene networks are parsimonious. Mol. Syst. Biol. 4:213. doi:10.1038/msb.2008.52.

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 42:425–432.

———. 1988b. Experimental studies of pleiotropy and epistasis in *Escherichia coli*. II. Compensation for maldaptive effects associated with resistance to virus T4. Evolution 42:433–440.

Lewis, M. 2009. An investigation into the evolution of relationships between species in an ecosystem. M.Sc. dissertation, ECS, University of Southampton, Southampton, U.K.

Lipson, H. 2007. Principles of modularity, regularity, and hierarchy for scalable systems. J. Biol. Phys. Chem. 7:125–128.

Lipson, H., J. B. Pollack, and N. P. Suh. 2002 On the origin of modular variation. Evolution. 56:1549–1556.

McGregor, S., V. Vasas, P. Husbands, and C. Fernando. 2012. Evolution of associative learning in chemical networks, PLoS Comput. Biol. 8:e1002739.

O'Reilly, R. C., and Y. Munakata. 2000. Computational explorations in cognitive neuro-science: understanding the mind by simulating the brain. MIT Press, Cambridge, MA.

Parter, M., N. Kashtan, and U. Alon. 2008. Facilitated variation: how evolution learns from past environments to generalize to new environments. PLoS Comput. Biol. 4:e1000206.

Pavlicev, M., and G. P. Wagner. 2012. A model of developmental evolution: selection, pleiotropy and compensation. Trends Ecol. Evol. 27:316–322.

Pavlicev, M., J. P. Kenney-Hunt, E. A. Norgard, C. C. Roseman, J. B. Wolf, and J. M. Cheverud. 2008. Genetic variation in pleiotropy: differential epistasis as a source of variation in the allometric relationship between long bone lengths and body weight. Evolution 62:199–213.

Pavlicev, M., J. M. Cheverud, G. P. Wagner. 2011. Evolution of adaptive phenotypic variation patterns by direct selection for evolvability Proc. R. Soc. B: Biol. Sci. 278:1903–1912.

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

Riedl, R. 1978. Order in living organisms. Wiley, New York.

Rumelhart, D. E., L. James, and McClelland and the PDP Research Group. 1986. Parallel distributed processing. Vols. I and II. MIT Press, Cambridge, MA.

Sansom, R. 2011. Ingenious genes: how gene regulation networks evolve to control development. MIT Press, Cambridge, MA.

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

Sniegowski, P. D., and H. A. Murphy. 2006. Evolvability. Curr. Biol. 16: R831–R834.

Stewart, J. 1997. The evolution of genetic cognition. J. Social Evol. Syst. 20:53–73.

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

Valiant, L. 2009. Evolvability. J. Assoc. Comput. Machinery 56:3:1–3:21.

Vohradsky, J. 2001a. Neural model of the genetic network. J. Biol. Chem. 276:36168–36173.

———. 2001b. Neural network model of gene expression. FASEB J. 15:846–854.

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

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

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

Wagner, G. P., and M. D. Laubichler. 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., M. Pavlicev, and J. M. Cheverud. 2007. The road to modularity. Nat. Rev. Genet. 8:921–931.

Watson, R. A., C. L. Buckley, R. Mills, and A. P. Davies. 2010a. Associative memory in gene regulation networks. Pp. 194–202 *in* H. Fellermann, M. Doerr, M. Hanczyc, L. L. Laursen, S. Maurer, D. Merkle, P.-A. Monnard, K. Støy, and S. Rasmussen, eds. Artificial life XII: proceedings of the twelfth international conference on the synthesis and simulation of living systems. MIT Press, Cambridge, MA.

Watson, R. A., R. Mills, C. L. Buckley, A. Penn, A. Davies, J. Noble, and S. Bullock. 2010b Adaptation without natural selection. Pp. 80–81 *in* Artificial life XII: proceedings of the twelfth International Conference on the synthesis and simulation of living systems. MIT Press, Cambridge, MA.

Watson, R. A., C. L. Buckley, and R. Mills. 2010c. Optimisation in 'self-modelling' complex adaptive systems. Complexity 16:17–26.

Watson, R. A., R. Mills, and C. L. Buckley. 2010d. Global adaptation in networks of selfish components: emergent associative memory at the system scale. Artif. Life 17:147–166.

Watson, R. A., R. Mills, and C. L. Buckley. 2011a. Transformations in the scale of behaviour and the global optimisation of constraints in adaptive networks. Adapt. Behav. 19:227–249.

Watson, R. A., D. M. Weinreich, and J. Wakeley. 2011b. Genome structure and the benefit of sex. Evolution 65:523–536.

Watson, R. A., N. Palmius, R. Mills, S. T. Powers, and A. S. Penn. 2011c. Can selfish symbioses effect higher-level selection? Pp. 27–36 *in* G. Kampis, I. Karsai, E. Szathmáry, eds. Advances in artificial life. Darwin Meets von Neumann. Springer, Berlin.

Watson, R. A., R. Mills R, C. L. Buckley, K. Kouvaris, A. Jackson, S. T. Powers, C. Cox, S. Tudge, P. Ryan, and 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., E. van Someren, and M. Reinders. 2001. A comparison of genetic network models. Pacific Sympos. Biocomput. 6:508–519.

West-Eberhard, M.-J. 2003. Developmental plasticity and evolution. Oxford Univ. Press, New York.

Wimsatt, W. C. 2001. Generative entrenchment and the developmental systems approach to evolutionary processes. Pp. 219–237 *in* S. Oyama, P. E. Griffiths, and R. D. Gray, eds. Cycles of contingency: developmental systems and evolution. MIT Press, Cambridge, MA.

Young, N. M., and B. Hallgrímsson. 2005. Serial homology and the evolution of mammalian limb covariation structure. Evolution, 59:2691–2704.

Young, N. M., G. P. Wagner, and B. Hallgrímsson. 2010. Development and the evolvability of human limbs. Proc. Natl. Acad. Sci. 107:3400–3405.

Associate Editor: J. Hunt

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

(a) Mean and relational effects of interaction terms.
(b) Hill-climbing model of selection.
(c) Effect of evolved interaction terms on the shape of phenotypic distributions.
(d) The cumulative effect of many small mutations under directional selection.
(e) Developmental switching–abrupt change in phenotype from linear change in genotype.
(f) Generalisation and evolvability.
(g) Modularly varying environments produces modular structure in the interaction network.
(h) Brief investigation of robustness to environmental noise.
(i) Robustness without reducing phenotypic variability.
(j) Recall from partial stimuli showing phenotypes over developmental time.

Supporting text (a)–(j) is provided including the following figures:
Figure S1. The effect of evolved regulatory interactions (after 10,000 generations) on the correlation of expression.
Figure S2. Quantified detail of Figure 3(f)—match of *G* and adult phenotypes to targets *S1* and *S2*.
Figure S3. Interaction coefficients evolved in modularly varying selective environment and resultant phenotypes.
Figure S4. Example phenotypes at developmental time step 10, after an environmental perturbation (randomization) to the developing gene expression pattern at successively later developmental time steps.
Figure S5. An evolved gene regulation network exhibits a developmental memory of phenotypes that have been selected for in the past.