PLOS Computational Biology

Phenotypes to remember: Evolutionary developmental memory capacity and robustness

--Manuscript Draft--

Manuscript Number:	PCOMPBIOL-D-20-00996
Full Title:	Phenotypes to remember: Evolutionary developmental memory capacity and robustness
Short Title:	Phenotypes to remember: Evolutionary developmental memory capacity and robustness
Article Type:	Research Article
Keywords:	gene regulatory network; associative neural networks, evolution; learning; evo-devo
Abstract:	There is increased awareness of the possibility of developmental memories resulting from evolutionary learning. Genetic regulatory and neural networks can be modelled by analogous formalism raising the important question of productive analogies in principles, processes and performance. We investigate the formation and persistence of various developmental memories of past phenotypes asking how the number of remembered past phenotypes scales with network size, to what extent memories stored form by Hebbian-like rules, and how robust these developmental "devo-engrams" are against networks perturbations (graceful degradation). The analogy between neural and genetic regulatory networks is not superficial in that it allows knowledge transfer between fields that used to be developed separately from each other. Known examples of spectacular phenotypic radiations could partly be accounted for in such terms.
Additional Information:	
Question	**Response**
Financial Disclosure Enter a financial disclosure statement that describes the sources of funding for the work included in this submission. Review the submission guidelines for detailed requirements. View published research articles from *PLOS Computational Biology* for specific examples. This statement is required for submission and **will appear in the published article** if the submission is accepted. Please make sure it is accurate.	YES A.S. and E.S. are supported by the National Research, Development and Innovation Office (NKFIH) under OTKA grant numbers K124438 (A.S.), K119347 (E.S. and A.S.) and GINOP-2.3.2-15-2016-00057 (E.S., P.S. and A.S.) research grants; A.S. was supported by the Bolyai János Research Fellowship of the Hungarian Academy of Sciences and the ÚNKP-19-4 New National Excellence Program of the Ministry of Human Capacities. M.S. is supported by the MTA Distinguished Guest Fellowship Programme in Hungary. All authors are supported by Templeton Foundation under grant number TWCF0268 ("Learning in evolution, evolution in learning").

Unfunded studies Enter: *The author(s) received no specific funding for this work.* **Funded studies** Enter a statement with the following details: • Initials of the authors who received each award • Grant numbers awarded to each author • The full name of each funder • URL of each funder website • Did the sponsors or funders play any role in the study design, data collection and analysis, decision to publish, or preparation of the manuscript? • **NO** - Include this sentence at the end of your statement: *The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.* • **YES** - Specify the role(s) played. * typeset	
Competing Interests Use the instructions below to enter a competing interest statement for this submission. On behalf of all authors, disclose any competing interests that could be perceived to bias this work—acknowledging all financial support and any other relevant financial or non-financial competing interests. This statement **will appear in the published article** if the submission is accepted. Please make sure it is accurate. View published research articles from *PLOS Computational Biology* for specific examples.	The authors have declared that no competing interests exist.

NO authors have competing interests Enter: *The authors have declared that no competing interests exist.* **Authors with competing interests** Enter competing interest details beginning with this statement: *I have read the journal's policy and the authors of this manuscript have the following competing interests: [insert competing interests here]* * typeset	
Data Availability Authors are required to make all data underlying the findings described fully available, without restriction, and from the time of publication. PLOS allows rare exceptions to address legal and ethical concerns. See the PLOS Data Policy and FAQ for detailed information. A Data Availability Statement describing where the data can be found is required at submission. Your answers to this question constitute the Data Availability Statement and **will be published in the article**, if accepted. **Important:** Stating 'data available on request from the author' is not sufficient. If your data are only available upon request, select 'No' for the first question and explain your exceptional situation in the text box. Do the authors confirm that all data underlying the findings described in their manuscript are fully available without restriction?	Yes - all data are fully available without restriction
Describe where the data may be found in full sentences. If you are copying our sample text, replace any instances of XXX with the appropriate details.	All relevant data are within the manuscript and its Supporting Information files.

• If the data are **held or will be held in a public repository**, include URLs, accession numbers or DOIs. If this information will only be available after acceptance, indicate this by ticking the box below. For example: *All XXX files are available from the XXX database (accession number(s) XXX, XXX.*). • If the data are all contained **within the manuscript and/or Supporting Information files**, enter the following: *All relevant data are within the manuscript and its Supporting Information files.* • If neither of these applies but you are able to provide **details of access elsewhere**, with or without limitations, please do so. For example: *Data cannot be shared publicly because of [XXX]. Data are available from the XXX Institutional Data Access / Ethics Committee (contact via XXX) for researchers who meet the criteria for access to confidential data.* *The data underlying the results presented in the study are available from (include the name of the third party and contact information or URL).* • This text is appropriate if the data are owned by a third party and authors do not have permission to share the data. * typeset	
Additional data availability information:	Tick here if the URLs/accession numbers/DOIs will be available only after acceptance of the manuscript for publication so that we can ensure their inclusion before publication.

Phenotypes to remember: Evolutionary developmental memory capacity and robustness

András Szilágyi[1,2,3], Péter Szabó[1,4], Mauro Santos[2,5] and Eörs Szathmáry[1,2,3*].

[1] Institute of Evolution, Centre for Ecological Research, 8237 Tihany, Hungary.

[2] Department of Plant Systematics, Ecology and Theoretical Biology, Eötvös Loránd University, 1117 Budapest, Hungary.

[3] Center for the Conceptual Foundations of Science, Parmenides Foundation, 82049 Pullach/Munich, Germany.

[4] Department of Ecology, Institute for Biology, University of Veterinary Medicine Budapest, 1078 Budapest, Hungary.

[5] Department de Genètica i de Microbiologia, Grup de Genòmica, Bioinformàtica i Biologia Evolutiva (GBBE), Universitat Autonòma de Barcelona, 08193 Bellaterra (Barcelona), Spain.

*Corresponding author

szathmary.eors@gmail.com (ES)

All authors contributed equally to this work.

Abstract

There is increased awareness of the possibility of developmental memories resulting from evolutionary learning. Genetic regulatory and neural networks can be modelled by analogous formalism raising the important question of productive analogies in principles, processes and performance. We investigate the formation and persistence of various developmental memories of past phenotypes asking how the number of remembered past phenotypes scales with network size, to what extent memories stored form by Hebbian-like rules, and how robust these developmental "devo-engrams" are against networks perturbations (graceful degradation). The analogy between neural and genetic regulatory networks is not superficial in that it allows knowledge transfer between fields that used to be developed separately from each other. Known examples of spectacular phenotypic radiations could partly be accounted for in such terms**.**

Introduction

Alan Turing, the father of machine learning, also formulated one of the most important mathematical models in developmental biology: the reaction-diffusion model for pattern generation [1]. This is striking because only recently a conceptual analogy between evolutionary developmental processes and artificial neural network-based learning models has been articulated [2]. Since development is the process whereby the phenotype is specified by the evolving genotype, late-evolved morphologies or functional capacities retain aspects of earlier stages ("memory") that were likely shaped by natural selection. These earlier stages might become reactivated if they are again useful in a different or a changing environment [3]. In this formulation evolutionary changes provide no novelty, defined as a structure that is non-homologous to an ancestral or existing one [4, 5], but allow for recursion. For instance, mimetic color patterns of an extinct morph of the butterfly *Heliconius cydno*, presumably as a result of human disturbance, can be reconstructed from wild-caught butterflies [6], meaning that the morph could recur in nature if the former conditions reappear. Also surprising is the repeatability of evolution among closely related lineages [7, 8]. An iconic textbook example is the extraordinary morphological convergence associated with adaptation to distinct ecological niches in cichlid fishes [9], with a large taxonomic diversity in the African Great Lakes Tanganyika (the oldest radiation, around 9-12 Myr ago with about 250 species), Malawi (less than 0.8 Myr ago and over 700 species) and Victoria (about 700 species evolved within the past 15,000 years) [10].

The idea that developmental processes can retain a memory of past selected phenotypes [11], together with the eerie ability of genomes to find adaptive solutions that quickly converge upon remarkably similar states ("attractors" [12]) in closely related lineages, clearly suggests a non-linear genotype-phenotype mapping capable of producing multiple distinct phenotypes [11,

13]. Non-linearity is also the hallmark of reaction-diffusion (Turing) and signaling systems involved in patterning processes [14], and developmental evolutionary biology (evo-devo) views the genotype-phenotype mapping as highly non-linear [15, 16]. Furthermore, it might not be farfetched to think of some sort of developmental memory in the cichlid adaptive radiation. The explosive diversification in Lake Victoria was predated by an ancient admixture between two distantly related riverine lineages, one from the Upper Congo and one from the Upper Nile drainage [17]. Many phenotypic traits known to contribute to the adaptation of different ecological niches in the Lake Victoria radiation are also divergent between the riverine species [18, 19]. Thus, when referring to the anatomical and morphological variation of Haplochromine cichlids, which are at the origin of the Lake Victoria radiation [17], Greenwood writes [19, p. 266]: "It is amongst the species of these various lacustrine flocks that one encounters the great range of anatomical, dental and morphological differentiation usually associated with the genus. The fluviatile species appear to be less diversified, but even here there is more diversity than is realized at first." If the high diversity in the Haplochromine cichlids of Lake Victoria is, to some extent, the result of re-evolved (similar) phenotypes in the ancestral fluviatile lineages, then the enduring question of why such an explosive diversification happened within a short time interval might have a simpler solution than previously thought. We aim here to sketch what the solution could be.

The genomic program for development operates primarily by the regulatory inputs and functional outputs of control genes that constitute network-like architectures [20], which are mathematically equivalent to artificial neural networks [21, 22]. Although the insights of Vohradsky [21, 22] and Watson et al. [11] shed light on an important analogy between neural and genetic regulatory networks, the conclusion of the theory of autoassociative networks cannot yet be readily extended to developmental systems. This is because of the different state space representations, as well as the nature of the task to be solved. Models of autoassociative networks

tend to work with positive/negative state variables (inherited from ferromagnetic systems). In contrast to this, in ontogenetic systems the relevant space is that of nonnegative real numbers, corresponding to concentrations of different molecules. Due to the nonlinear activation function features of models working with the above mentioned, alternative state representations can markedly differ. Another important consideration is that autoassociative networks (as their name indicates) solve the problem of the recovery of a particular state (attractor property). During ontogeny we require something more: not only should the adult stage be stable, but the system should reach this state from a particular embryo state (heteroassociative property). It can easily happen that one has to find a solution to a problem of the transition from multiple embryo states and the corresponding adult states. In short, we require networks that solve problems of auto/heteroassociativity in one, as we model here.

Methods

Developmental model

Our model is a formal description of ontogenetic development operating primarily by the regulatory inputs and functional outputs of control genes. Consider an organism with N genes. Its developmental state at time *t*, expressed by its gene product composition (e.g., proteins), can be represented by the vector $\mathbf{p}(t) = (p_1, p_2, \dots, p_N)^{\mathrm{T}}$ with each element being the quantity of the product of a gene. These quantities are assumed to change due to protein decay and gene expression processes. Following [11], the ontogenetic dynamics of the developmental state can be described by the difference equation

$$p_i(t+1) = (1-\delta)p_i(t) + \tau f([\mathbf{M}\mathbf{p}(t)]_i), \qquad (1)$$

where τ denotes the decay rate, τ denotes the maximal gene expression rate, $f(.)$ is the activation function, and the matrix $\mathbf{M}$ stands for the regulatory network. An m_{ij} entry of the matrix gives the regulatory effect of the product of gene j on the expression level of gene i; positive and negative elements imply activation and inhibition, respectively. The cumulative regulatory effects on any single gene i, i.e. the ith element of the product $\mathbf{Mp}$, determine the gene expressions via a sigmoid activation function modelled here as $f(x) = (1 + \tanh(\omega x))/2$, where ω is the slope parameter.

From an ontogenetic viewpoint, the task of the gene regulatory network is to guide the individual along a developmental pathway from an initial embryonic state $\mathbf{p}(0) = \mathbf{e}$ to a specific adult state $\mathbf{p}(T) \rightarrow \mathbf{a}$. In real systems, an ensemble of different developmental pathways is desired, each responsible for achieving some environment-specific adult state from a particular embryonic state. We used T = 150 iteration to reach the steady state.

Evolutionary model

In the evolutionary model we considered a population of K individuals, with each member of the population represented by its regulatory matrix. All the interaction matrix elements were zero initially, representing an undeveloped regulation. Every individual shared the same environment. We assumed $Q = 3$ number of different selective environments, each defining an embryonic state $\mathbf{e}^{(q)}$ and a corresponding adapted adult state $\mathbf{a}^{(q)}$. The selective environments alternated randomly; if the average fitness of the population approached the optimum ($\bar{w} > 0.95$ for at least 20 consecutive generations), or after 10000 generations, a new environment was chosen at random. In each generation the individuals underwent mutation, development and selection steps as follows.

Mutation*:* The mutation of the regulation network was implemented by adding a normally distributed random value, with zero mean and μ_W variance, to a randomly selected matrix element. Matrix elements were clipped into the range $[-1,1]$.

Development*:* We got the equilibrium, adult state of each member of the population by iterating Eq. (1).

Selection*:* The fitness of individual *k* was expressed by a similarity index derived from the Euclidean distance between the actual adult state $\mathbf{p}(T)$ and the environment-specific optimal adult phenotype $\mathbf{a}^{(q)}$ as

$$w_k = 1 - \sqrt{\sum_{n=1}^{N} \left[\frac{p_n(T) - a_n^{(q)}}{\tau/\delta} \right]^2}. \qquad (2)$$

Then the regulatory matrix of a randomly selected individual was replaced by that of the individual with the highest fitness (elitist selection).

Embryonic and (optimally adapted) adult vectors*:* The number of genes was $N = 100$ with a low average expression level of $\sigma = 0.1$, where 40% of the expressed genes were common, 20% were partially common, and 40% were unique in all the embryonic and all the adult vectors. Specifically, the expression sites of the employed state vectors were

$\mathbf{e_1} = \{\underline{13}, \underline{19}, \underline{32}, \overline{36}, 39, 49, 55, \underline{72}, 81, \overline{87}\}$, $\mathbf{e_2} = \{\underline{13}, \underline{19}, 31, \underline{32}, \overline{40}, 60, 62, \underline{72}, \overline{87}, 100\}$,

$\mathbf{e_3} = \{5, \underline{13}, \underline{19}, \underline{32}, \overline{36}, \overline{40}, 47, 67, \underline{72}, 94\}$, $\mathbf{a_1} = \{\underline{6}, \overline{12}, 20, 24, \underline{46}, 65, \underline{84}, \overline{86}, 88, \underline{92}\}$,

$\mathbf{a_2} = \{\underline{6}, 11, 28, \underline{46}, 79, \underline{84}, \overline{86}, \overline{91}, \underline{92}, 96\}$, $\mathbf{a_3} = \{\underline{6}, \overline{12}, \underline{46}, 56, 61, 66, 80, \underline{84}, \overline{91}, \underline{92}\}$; where underlines and overlines denote the common and partially common elements, respectively. The initial state was always a perturbed embryonic state. The perturbation was performed, similar to the mutations, by adding a normally distributed random value, with zero mean and μ_e variance, to

a randomly selected element of the environment-specific embryo vector. Vector elements were clipped into the range $[0, \tau/\delta]$.

Perturbation analysis

The embryo states were perturbed by flipping the vector elements from low to high, or vice versa, with the given probability. The interaction matrices were perturbed by either adding random values to all matrix elements, drawn from a normal distribution with the given standard deviation, or by nullifying a proportion of the elements. Matrix elements were clipped into the range $[-1, 1]$. Note that, in the evolutionary algorithm we perturbed only single elements of the embryonic states. In contrast, in the analytical matrix construction we perturbed all elements of the embryonic vectors to incorporate the accumulating effects of many consecutive perturbations on the interaction matrix.

Results

To perform the developmental task, the network must guarantee that (i) each adult state is a stable equilibrium point of the dynamics (stability condition), and (ii) each embryonic state is within the basin of attraction of its corresponding adult state (attraction condition); these two conditions correspond to the auto- and heteroassociative properties in a neural network [23]. Note that this is a more difficult task than a simple pattern recovery problem, which is known to be achievable by a neural network with the standard Hebbian learning rule that fulfils only the stability condition [24]. Not only must all the adult states have a basin of attraction, but these basins must include the corresponding embryonic states.

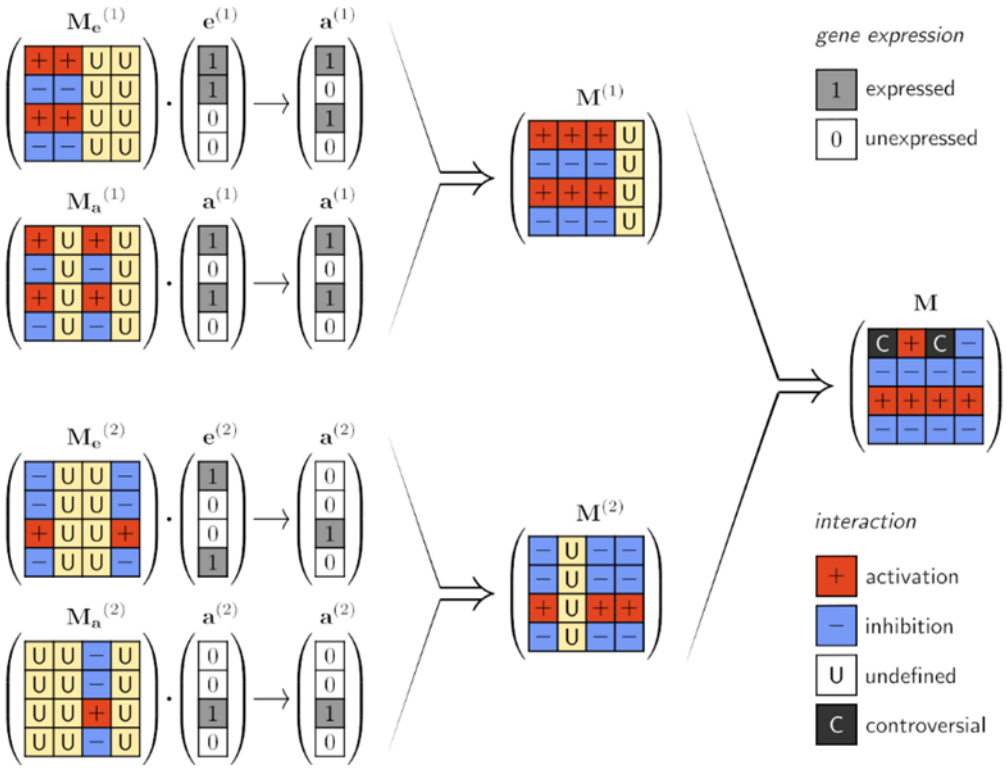

Fig. 1. Illustration of the construction rules of interaction matrices based on theoretical considerations on the optimal pairwise interaction types between genes. $\mathbf{e}^{(1)}$ and $\mathbf{a}^{(1)}$ are the first embryo-adult pair, $\mathbf{e}^{(2)}$ and $\mathbf{a}^{(2)}$ the second pair. Depending on the combination of gene expressions $e_i^{(n)}$ and $a_i^{(n)}$ in an embryo-adult vector pair ($n = 1,2$), an m_{ij} element of the interaction matrix can be positive ($'+'$, activation), negative ($'-'$, inhibition), or undefined ($'\mathrm{U}'$). To ensure correct development $\left(\mathbf{M}_\mathbf{e}^{(n)}\mathbf{e}^{(n)} \rightarrow \mathbf{a}^{(n)}\right)$ the $\mathbf{M}_\mathbf{e}^{(n)}$ matrices must have the structure indicated in the figure. (If $e_j^{(n)} = 1$ and $a_i^{(n)} = 1$, then $m_{ij}^{(n)} = '+'$; if $e_j^{(n)} = 1$ and $a_i^{(n)} = 0$, then $m_{ij}^{(n)} = '-'$; if $e_j^{(n)} = 0$, then $m_{ij}^{(n)} = '\mathrm{U}'$; irrespective of the value of $a_i^{(n)}$.) A similar argument holds for the stability criteria $\left(\mathbf{M}_\mathbf{a}^{(n)}\mathbf{a}^{(n)} \rightarrow \mathbf{a}^{(n)}\right)$ and results in the $\mathbf{M}_\mathbf{a}^{(n)}$ matrices. By combining $\mathbf{M}_\mathbf{e}^{(n)}$ and $\mathbf{M}_\mathbf{a}^{(n)}$ the resulting $\mathbf{M}^{(n)}$ fulfills both the attractivity and stability criteria. The combination rules are the following: $(+,+) \rightarrow +$; $(-,-) \rightarrow -$; $(\pm, \mathrm{U}) \rightarrow \pm$; and $(\pm, \mp) \rightarrow \mathrm{C}$, which can be done practically by taking the element-wise average of the two matrices. The ultimate combination of all $\mathbf{M}^{(n)}$s results in a matrix that fulfills the attraction and stability criteria for all different embryo-adult pairs.

We found that the task-optimized structure of the regulatory network can be inferred from the embryo-adult state vector pairs in the form of an interaction matrix **M** (Fig 1). Consider the simplest case with one embryo-adult pair (i.e. one developmental pathway). Depending on whether

a gene is expressed in the adult state or not, all the other expressed gene products, in either the embryonic or the adult state, must enhance or block its expression, respectively. This would provide, on the one hand, stability for the adult state and, on the other hand, attraction from the embryonic state. Note, however, that if a gene is expressed in neither the embryonic nor the adult state, then its regulatory effect is irrelevant, therefore the corresponding matrix elements are undetermined. In summary, an m_{ij} element of the regulatory matrix **M** should be positive or negative, depending on whether the *i*th gene is expressed in the adult state or not, except when the *j*th gene is expressed in neither the embryonic nor the adult state. The above line of thought can be generalized for arbitrary Q number of embryo-adult state pairs. Denoting the zero-one normalized embryonic and adult state vectors by **e** and **a**, such a matrix can be obtained by averaging two dyadic products for all developmental pathways as

$$\mathbf{M} = \frac{1}{2Q}\sum_{q=1}^{Q}\left(2\hat{\mathbf{a}}^{(q)} - 1\right) \circ \hat{\mathbf{a}}^{(q)} + \left(2\hat{\mathbf{a}}^{(q)} - 1\right) \circ \hat{\mathbf{e}}^{(q)}, \qquad (3)$$

where Q stands for the number of embryo-adult state pairs and q denotes the different pairs. The first and second dyadic products are responsible for the stability and attraction conditions, respectively. Within each dyadic product the right argument determines whether an entry is relevant from the viewpoint of the state vector, whereas the left dyadic argument determines its sign. The resulting matrix contains positive values, negative values and zeros for activator, inhibitory and undetermined elements, respectively. Notice that the developmental pathways can be in conflict with each other as to whether a gene should be up- or downregulated by another gene. It is instructive to compare this formula with the standard Hebbian learning rule $\mathbf{H} = \mathbf{a} \circ \mathbf{a}$ for $a_i \in \{-1, +1\}$. Its modification for $a_i \in \{0, +1\}$ vectors that preserves that stability condition is $\mathbf{H} = (2\mathbf{a} - 1) \circ \mathbf{a}$, which is identical to the first term in Eq. (3), c.f. Table 1.

Table 1. Comparison of the resulting interaction matrices for an autoassociative task with the two representations.

	$\{-1,+1\}$ **representation**	$\{0,+1\}$ **representation**
learning rule	$H_{ij}=\sum_q a_i^{(q)}a_j^{(q)}$ (Hebb-rule)	$M_{ij}=\sum_q (2a_i^{(q)}-1)a_j^{(q)}$
symmetry	$H_{ij}=H_{ji}$	non-symmetric
neutralities in weight matrix	no neutral elements	can be neutral elements ("opposite to" zero vector elements)
main diagonal	always positive (if allowed)	can be negative or positive
robustness against perturbation of matrix elements	smaller, due to the absence of neutral elements	larger, due to the presence of neutral elements
structure	has a unique structure	many different realizations

We investigated the parameter dependence of the analytic model. As for the regulatory matrix we used a slightly modified version of Eq. (3). Treves [25] claims that the interaction terms should be modified by the average expression σ, i.e. the proportion of expressed genes. This is because if a larger proportion of genes is expressed, then proportionally smaller interaction strengths are needed for the same regulatory effect on any single gene. Incorporating this consideration into Eq. (3) gives

$$\mathbf{M}=\frac{1}{2Q}\sum_{q=1}^{Q}\left(2\hat{\mathbf{a}}^{(q)}-1\right)\circ\left(\hat{\mathbf{a}}^{(q)}-\sigma+\hat{\mathbf{e}}^{(q)}-\sigma\right). \qquad (4)$$

The performance of a regulatory network constructed by the above rule changes with the number of developmental pathways and gene expression levels (Fig 2). With increasing number of embryo-adult pairs, the accumulating conflicts between them inevitably corrupt the regulatory ability of the network; some adult states will be unreachable from their embryonic states. Nevertheless, the network is able to tolerate a fair number of conflicts, related to its structural stability. Since conflicts can occur only between non-orthogonal state vectors, the performance of the network also depends

on the amount of overlap in the expression patterns of states belonging to different pairs. This highlights the importance of the proportion of expressed genes; i.e., the sparseness of the state vectors. If these vectors are very sparse, then they are unlikely to be non-orthogonal, therefore the number of learnable embryo-adult pairs is quite large. The number of learnable pairs is a decreasing function of sparseness, because the numerous non-orthogonal state vectors and the largely different gene expressions in the adult states lead to several conflicts among them. Regarding the effect of system size on functionality, the results are in line with the expectations; the higher the number of genes, the higher is the number of "error free" developmental pathways (Fig 3).

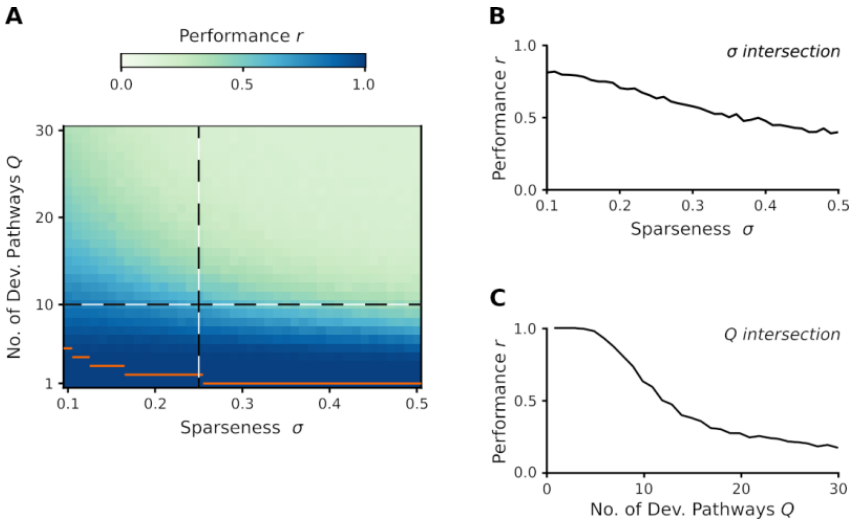

Fig. 2. Performance of the analytic developmental networks. We assumed different sparseness (proportion on non-zero entries in the state vectors) values and different number of embryo-adult pairs. Embryo and desired adult vectors were generated by independently setting each vector element to high or low randomly according to the sparseness value.The performance was measured by the averaged (over 400 realizations) Pearson correlation(s) between the desired and the experienced adult state(s) for all developmental pathways (panel A). Panels B and C show a more detailed view for the two cross-sections of the parameter space (indicated by dashed lines in panel A). Orange horizontal lines show the maximum number of orthogonal state vectors for the given sparseness values. Parameters: $N = 100, \delta = 0.2, \tau = 1, \omega = 25$).

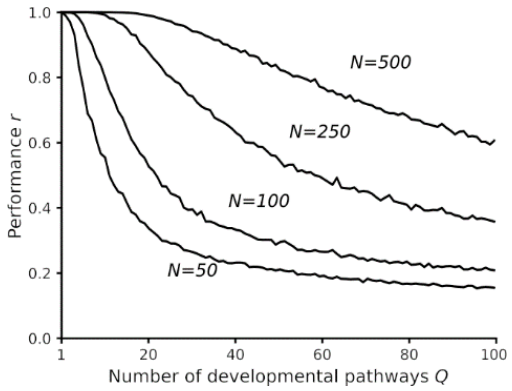

Fig. 3. Performance of the analytic developmental networks with different number of genes and developmental pathways. Relevant parameters as in Fig. 2 and $\sigma = 0.1$.

A key question is whether a functional network is attainable by Darwinian selection via a series of mutation-selection steps. In our evolutionary model we used a more realistic Darwinian dynamics than the solitary stochastic hill climbing [11]. From the viewpoint of the theory of artificial neural networks this process can be regarded as a Darwinian dynamics-driven learning process. The evolutionary algorithm yields interaction matrices that contain positive and negative values where the heuristic formulation predicts them (Fig 4). While the individual interaction matrices vary, their average is in line with the heuristically derived matrix. The values are arranged into a characteristic structure; positive and negative entries form horizontal stripes, intermitted with vertical stripes of near-zero values (c.f. Fig 1). Those genes have the largest effect on the developmental process, which are expressed in any embryonic or adult states (c.f. marked columns in Fig 4). Depending on whether the affected gene is expressed in any of the adult states, they have a strong positive or negative effect (c.f. marked rows in Fig 4). The rest of the genes drift freely in individual realizations due to a lack of selective pressure. Consequently, the average values in these positions are approximately zero (c.f. grey columns in left panel of Fig 4). The corresponding values in the analytic treatment (undefined elements) are zero by definition. The only major difference from the heuristic matrix is that the main diagonal elements of the evolutionary matrix

are mainly negative, which means that the expression of every gene is under negative feedback by its own inhibitory product. A possible explanation is that without a strong negative feedback a gene could be easily overexpressed due to the perturbations of the interaction elements. This is more probable if the sparseness of the expression vectors is low, as it was in our case. This picture is likely to change with hierarchical developmental regulation, the evolution of which takes longer time and should be investigated in the future.

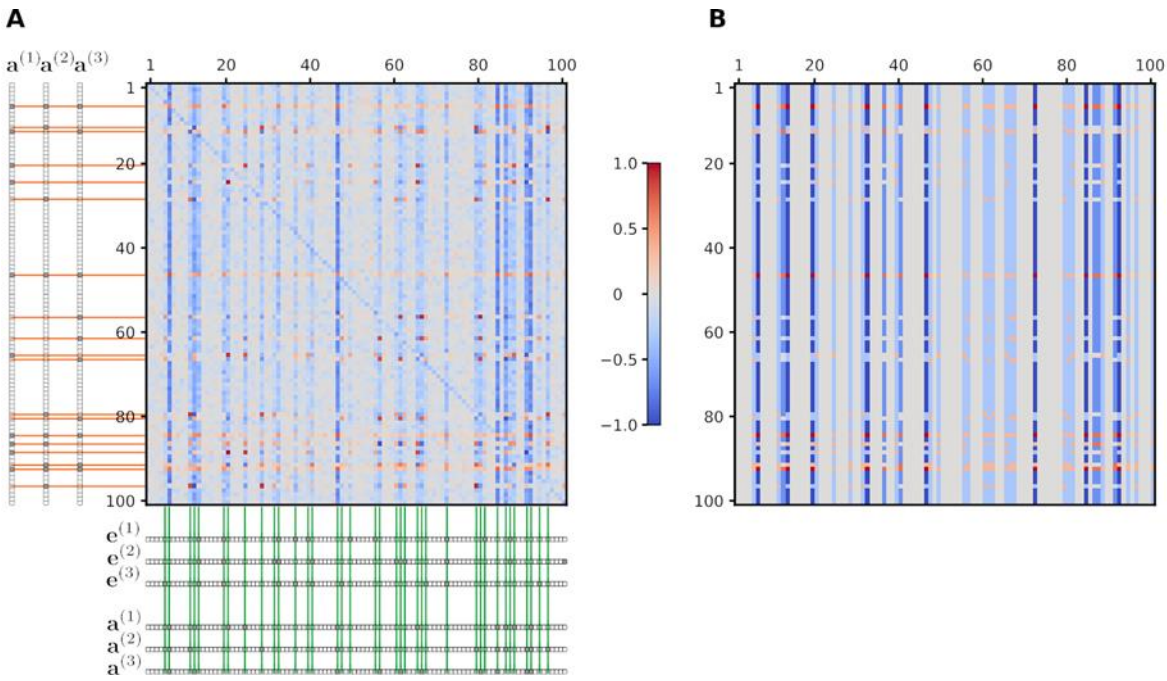

Fig. 4. Structure of the interaction matrix obtained with the evolutionary algorithm as compared to the analytically derived one for three developmental pathways. (A) The evolutionary interaction matrix was obtained by averaging the output of 300 independent runs of the evolutionary algorithm. The three applied environment-specific embryonic and adult state vectors are shown along the sides. Orange guidelines highlight those rows where the corresponding genes are expressed in at least one adult state, whereas green guidelines highlight those columns where at least one gene is expressed in any of the embryonic or adult states. (B) The theoretically predicted interaction matrix was constructed from the embryonic and adult state vectors using Eq. (3). Parameters as in Fig. 2 and $Q = 3, K = 100, \mu_W = 0.05, \mu_e = 0.1,\ \sigma = 0.1$.

A detailed view of the evolutionary process is shown in Fig 5. During the early generations, where the gene regulation is undeveloped, it takes many generations (i.e., mutation-selection steps) to approach the environment-specific optimum. In addition, selection for one environment can have adverse effects on performance in another environment if the basin of attraction of the actually selected adult state engulfs the neighborhood of the embryonic states of other adult states. But those interactions which are not beneficial in all environments are eliminated. Detrimental mutations may arise any time also in a well-functioning system, but selection eliminates them over the timescale of a few environmental changes.

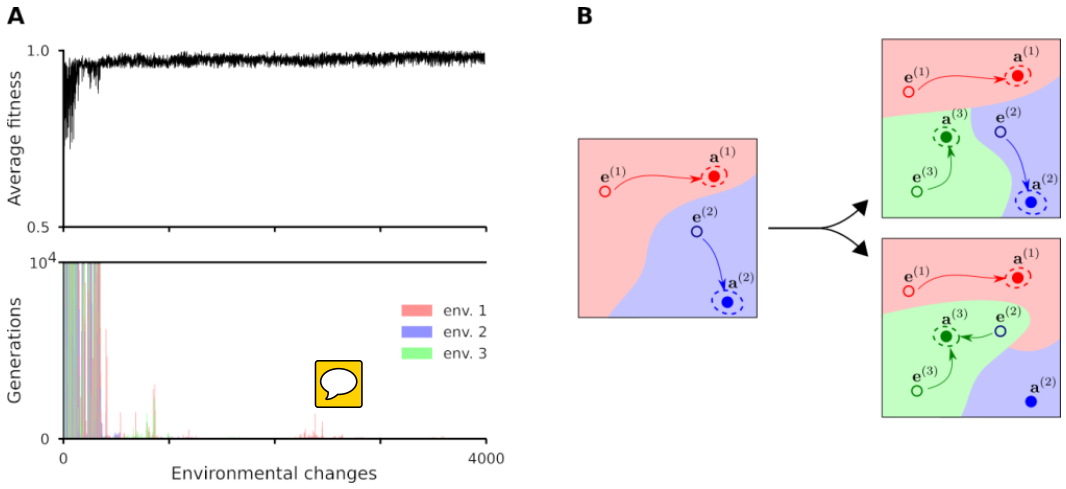

Fig. 5. Learning of three different developmental pathways in the evolutionary model. (A) Average fitness and the mutation-selection steps needed to achieve a well-functioning developmental network during random environmental changes. The three environments are denoted by red, green and blue. Parameters as in Fig. 4. (B) Schematic illustration of the changes in the state-space dynamics during the evolutionary process with three developmental pathways (indicated by red, blue and green colors). The panels show the basins of attraction of an initial, random regulation system with two embryo-adult pairs (left), a well-functioning one (top right) and a bad one, where the basins of attraction of the adult states (filled dots encircled by dotted lines indicating variation around the target phenotypes) include not only their corresponding embryo states (bottom right).

A developmental process must be sufficiently robust against stochastic perturbations of both the embryonic state and the gene interaction matrix. It requires that the neighborhood (according to a given metric) of the embryonic states must also be in the basin of their corresponding adult states. Therefore, some inputs of variation should produce little or no phenotypic variation at all, a phenomenon that has received a lot of attention under the labels of canalization, robustness or buffering [26-28]. The recovery performance of the network changes with increasing amount of perturbations (Fig 6). The system is very robust against perturbations regarding the embryonic state, and it is moderately robust against both additive perturbations and eliminated interactions regarding the interaction network. This robustness is attributable to the high number of neutral elements (correspond to the zero values in state vectors) of the interaction matrix. This is in sharp contrast to the standard Hebbian set-up, where there are no neutralities, due to the $\{-1, +1\}$ representation; see Table 1. Resilience understandably decreases with the number of developmental pathways in all cases, but conforming to "graceful degradation" in artificial neural networks; i.e., performance first decreases mildly and drops fast only beyond a critical strength of perturbation [24]. To sum up, variation is apportioned into discontinuous (basins of attraction) and continuous (small perturbations around the target) phenotypes (Fig 5B). Evo-devo mainly focus on the first kind of variation whereas standard evolutionary genetics focus on the second [15, 29].

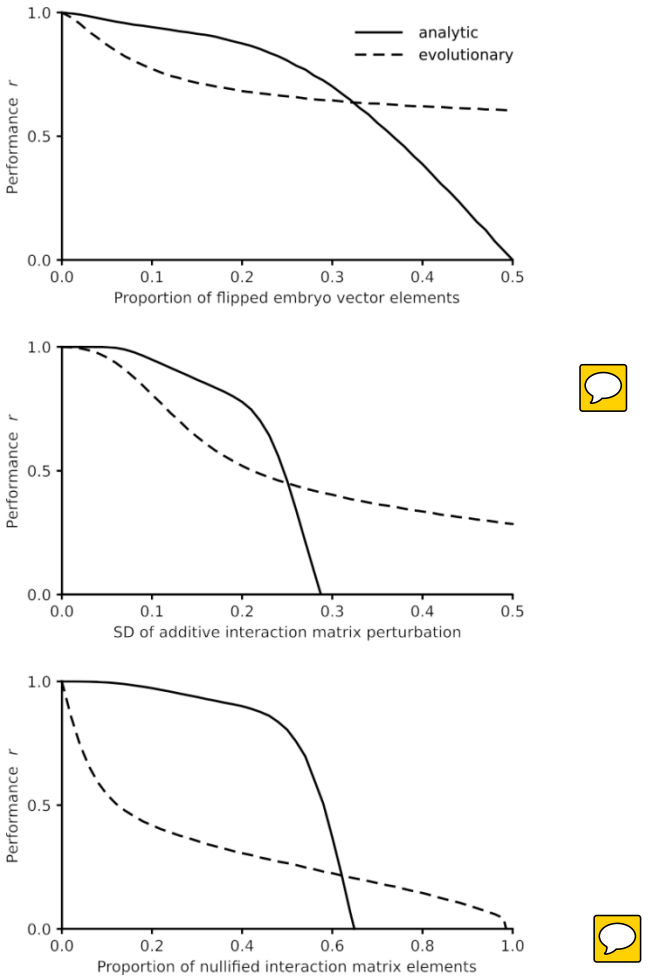

Fig. 6. Robustness of the developmental dynamics against perturbations. The interaction matrices were constructed from the given number of embryo-adult vector pairs according to Eq. (4). The performance was expressed by the Pearson correlation(s) between the desired and the experienced adult state(s) for all developmental pathways after T=150 iterations averaged over 300 matrices and 100 perturbations for each parameter combination. (A**)** Performance against the proportion of the flipped embryonic vector elements. (B**)**. Performance against the standard deviation (SD) of the perturbation of the interaction matrix. All elements of the matrix were perturbed additively by an $N(0, SD)$ random number. (C**)** Performance against the proportion of nullified elements of the interaction matrix. Each element of the interaction matrix was set to zero with the given probability. Relevant parameters are as in Fig. 4.

Discussion

Treating gene regulatory networks as formally analogous to artificial neural networks [21, 22] allows translating the well-known dynamics of the latter [30] to model genomic programs for development. There is widespread natural variation in morphogenic pathways [31], and the developmental memory of past selected phenotypes [11] is akin to the memory capacity of neural networks. This developmental memory allows populations to re-evolve phenotypes much faster than it would be possible if they had to evolve de novo. Previous speculation on the role of the heat-shock protein Hsp90 as a capacitor for releasing hidden morphogenetic variation that could allow fast morphological radiations [31] has been criticized on the grounds that the function of Hsp90 is to prevent morphological aberrations. Furthermore, some sense of purposive evolution, fully incompatible with the lack of foresight of natural selection, lays behind this sort of interpretations [32].

These criticisms do not apply here because in our developmental model past selected states can recur in the population if they appear useful again in a different environment or body context. As any theoretical model, ours obviously has inherent limitations and highly simplifies the representation of biological systems. However, to the extent that it captures sufficient conditions to generate the phenomenon of morphological radiations, more complex explanations are not required. Thus, the assumption that structural novelties (or "key innovations") are associated with adaptive radiations into new ecological niches (e.g. [33, p. 159]) might be unwarranted. There is a noteworthy implication in the foregoing consideration for the understanding of atavism. Crocodilian teeth can grow in mutant birds, which suggests the reactivation of the associated developmental machinery [34], that required the resurrection [35] of a key aspect of regulation. The same neurons participate in the storage of different engrams in neural networks. The same

holds for the storage of devo-engrams in genetic regulatory networks. Resurrection leading to atavism requires only limited reactivation of a few connections in a network that is maintained by the current selective forces. An exciting question is how evo-devo learning can generalize from the "training set" (previously selected target phenotypes) to novel ones [11, 36]. Prediction [36] is that generalization potential works within a set that can be characterized by the same formal grammar.

While the theory of neural networks can (and does) infer the same conclusions based on different representations, in the case of modelling real biological situations the adequacy of the representation can be crucial (the same holds for neuronal networks). Our results show that a linear change to the representation has profound impact on the essential features of the system. While in the customary (neural) {-1,+1} representation there are no neutral elements in the interaction matrix, the biologically adequate {0,+1} representation of genetic regulatory networks allows for the free choice of interaction elements being opposite to "0". This feature turns out to increase the robustness of the system against the disturbance of interaction coefficients. Another feature of our representation is the large number of different interaction matrices entailing the same developmental process, thus evolution "from scratch" does not face so many constraints.

Acknowledgements

The authors thank Ferenc Jordán and Harold. P. de Vladar for comments on an earlier draft. A.S. and E.S. are supported by the National Research, Development and Innovation Office (NKFIH) under OTKA grant numbers K124438 (A.S.), K119347 (E.S. and A.S.) and GINOP-2.3.2-15-2016-00057 (E.S., P.S. and A.S.) research grants; A.S. was supported by the Bolyai János Research Fellowship of the Hungarian Academy of Sciences and the ÚNKP-19-4 New National Excellence Program of the Ministry of Human Capacities. M.S. is supported by the MTA Distinguished Guest Fellowship Programme in Hungary. All authors are supported by Templeton Foundation under grant number TWCF0268 ("Learning in evolution, evolution in learning"). The authors declare no competing interests. All the computer files employed in this study will be available upon acceptance.

References

1. Turing AM. The chemical basis of morphogenesis. Philos Trans R Soc B. 1952; 237:37-72.
2. Watson RA, Szathmáry E. How Can Evolution Learn? Trends Ecol Evol. 2016; 31:147-157.
3. Sharov AA. Evolutionary constraints or opportunities? Biosystems. 2014; 123:9-18.
4. Müller GB, Wagner GP. Novelty in evolution: Restructuring the concept. Annu Rev Ecol Syst. 1991; 22:229-256.
5. Peterson T, Müller GB. What is evolutionary novelty? Process versus character based definitions. J Exp Zool (Mol Dev Evol). 2013; 320B:345-350.
6. Linares M. The ghost of mimicry past: laboratory reconstitution of an extinct butterfly 'race'. Heredity. 1997; 78:628-635.
7. Orgogozo V. Replaying the tape of life in the twenty-first century. Interface Focus. 2015; 5:20150057.
8. Blount ZD, Lenski RE, Lossos JB. Contingency and determinism in evolution: Replying life's tape. Science. 2018; 362: eaam5979.
9. Muschick M, Indermaur A, Salzburger W. Convergent evolution within an adaptive radiation of cichlid fishes. Curr Biol. 2012; 22:2362-2368.
10. Salzburger W. Understanding explosive diversification through cichlid fish genomics. Nat Rev Genet. 2018; 19:705-717.
11. Watson RA, Wagner GP, Pavlicev M, Weinreich DM, Mills R. The evolution of phenotypic correlations and "developmental memory". Evolution. 2014; 68:1124-1138 (2014).
12. Conway Morris S. Life's solution. Inevitable humans in a lonely universe. Cambridge university press, 2003.
13. Kouvaris K, Clune J, Kounios L, Brede M, Watson RA. How evolution learns to generalise: Using the principles of learning theory to understand the evolution of developmental organisation. PLoS Comput Biol. 2017; 13:e1005358.
14. Murray JD. Mathematical Biology, 2nd ed. Springer, 1993.
15. Alberch P. Developmental constraints in evolutionary processes. In: Evolution and development. Dahlem Konferenzen 1982, Bonner JT, ed. Springer-Verlag, 1982, pp 313-332.

16. Gjuvsland AB, Vik JO, Beard DA, Hunter PJ, Omholt SW. Bridging the genotype–phenotype gap: what does it take? J Physiol. 2013; 591:2055-2066.
17. Meier JI, Marques DA, Mwaiko S, Wagner CE, Excoffier L, Seehausen O. Ancient hybridization fuels rapid cichlid fish adaptive radiations. Nat Commun. 2017; 8:14363.
18. Seehausen O, Koetsier E, Schneider MV, Chapman LJ, Chapman CA, et al. Nuclear markers reveal unexpected genetic variation and a Congolese-Nilotic origin of the Lake Victoria cichlid species flock. Proc R Soc Lond B Biol Sci. 2003; 270:129-137.
19. Greenwood PH. Towards a phyletic classification of the 'genus' Haplochromis (Pisces, Cichlidae) and related taxa. Part 1. Bull Br Mus Nat Hist Zool. 1979; 35:265-322.
20. Levine M, Davidson EH. Gene regulatory networks for development. Proc Natl Acad Sci USA. 2005; 102:4936-4942.
21. Vohradsky J. Neural model of the genetic network. J Biol Chem. 2001; 276:36168-36173.
22. Vohradsky J. Neural network model of gene expression. FASEB J. 2001; 15:846-854.
23. Hirahara M. "Associative memory" in Encyclopedia of Neuroscience. Binder MD, Hirokawa N, Windhorst U, eds. Springer, 2009, pp 195-201.
24. Rolls ET. Memory, attention and decision-making. Oxford university press, 2008.
25. Treves A. Graded-response neurons and information encodings in autoassociative memories. Phys Rev A. 1990; 42:2418-2430.
26. Waddington CH. Canalization of development and the inheritance of acquired characters. Nature. 1942; 150:563-565.
27. Nijhout HF. The nature of robustness in development. Bioessays. 2002; 24:553-563.
28. Hallgrimssona B, Greena RM, Katz DC, Fish JL, Bernier FP, et al. The developmental-genetics of canalization. Sem Cell Dev Biol. 2019; 88:67-79.
29. Müller GB. Evo-devo: extending the evolutionary synthesis. Nat Rev Genet. 2007; 8:943-949.
30. Haykin S. Neural networks: a comprehensive foundation, 2nd ed. Pearson Education, 1999.
31. Rutherford SL, Lindquist S. Hsp90 as a capacitor for morphological evolution. Nature. 1998; 396:336-342.
32. Dickinson WJ, Seger J. Cause and effect in evolution. Nature. 1999; 399:30.
33. Carroll SB, Grenier JK, Weatherbee SD. From DNA to diversity. Molecular genetics and the evolution of animal design, 2nd ed. Blackwell, 2005.

34. Harris MP, Hasso SM, Ferguson MWJ, Fallon JF. The development of Archosaurian first-generation teeth in a chicken mutant. Curr Biol. 2006; 16:371-377.
35. Marshall CR, Raff EC, Raff RA. Dollo's law and the death and resurrection of genes. Proc Natl Acad Sci USA. 1994; 91:12283-12287.
36. Parter M, Kashtan N, Alon U. Facilitated variation: how evolution learns from past environments to generalize to new environments. PLoS Comput Biol. 2008; 4:e1000206.