

RESEARCH ARTICLE

How to fit in: The learning principles of cell differentiation

Miguel Brun-Usan , Christoph Thies , Richard A. Watson*

Institute for Life Sciences/Electronics and Computer Sciences, University of Southampton, Southampton, (United Kingdom)

* R.A.Watson@soton.ac.uk

Abstract

Cell differentiation in multicellular organisms requires cells to respond to complex combinations of extracellular cues, such as morphogen concentrations. Some models of phenotypic plasticity conceptualise the response as a relatively simple function of a single environmental cue (e.g. a linear function of one cue), which facilitates rigorous analysis. Conversely, more mechanistic models such as those implementing GRNs allow for a more general class of response functions but make analysis more difficult. Therefore, a general theory describing how cells integrate multi-dimensional signals is lacking. In this work, we propose a theoretical framework for understanding the relationships between environmental cues (inputs) and phenotypic responses (outputs) underlying cell plasticity. We describe the relationship between environment and cell phenotype using logical functions, making the evolution of cell plasticity equivalent to a simple categorisation learning task. This abstraction allows us to apply principles derived from learning theory to understand the evolution of multi-dimensional plasticity. Our results show that natural selection is capable of discovering adaptive forms of cell plasticity associated with complex logical functions. However, developmental dynamics cause simpler functions to evolve more readily than complex ones. By using conceptual tools derived from learning theory we show that this developmental bias can be interpreted as a learning bias in the acquisition of plasticity functions. Because of that bias, the evolution of plasticity enables cells, under some circumstances, to display appropriate plastic responses to environmental conditions that they have not experienced in their evolutionary past. This is possible when the selective environment mirrors the bias of the developmental dynamics favouring the acquisition of simple plasticity functions—an example of the necessary conditions for generalisation in learning systems. These results illustrate the functional parallels between learning in neural networks and the action of natural selection on environmentally sensitive gene regulatory networks. This offers a theoretical framework for the evolution of plastic responses that integrate information from multiple cues, a phenomenon that underpins the evolution of multicellularity and developmental robustness.

OPEN ACCESS

Citation: Brun-Usan M, Thies C, Watson RA (2020) How to fit in: The learning principles of cell differentiation. *PLoS Comput Biol* 16(4): e1006811. <https://doi.org/10.1371/journal.pcbi.1006811>

Editor: Qing Nie, University of California Irvine, UNITED STATES

Received: January 24, 2019

Accepted: February 20, 2020

Published: April 13, 2020

Copyright: © 2020 Brun-Usan et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data are within the manuscript and its Supporting Information files.

Funding: This work was funded by the Templeton Foundation (Grant 66501 to MBU, ChT and RAW) www.templeton.org/grant/putting-the-extended-evolutionary-synthesis-to-the-test. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

Author summary

In organisms composed of many cell types, the differentiation of cells relies on their ability to respond to complex extracellular cues, such as morphogen concentrations, a phenomenon known as cell plasticity. Although cell plasticity plays a crucial role in development and evolution, it is not clear how, and if, cell plasticity can enhance adaptation to a novel environment and/or facilitate robust developmental processes. In some models, the relationships between the environmental cues (inputs) and the phenotypic responses (outputs) are conceptualised as one-to-one (i.e. simple 'reaction norms'); whereas the phenotype of plastic cells commonly depends on several simultaneous inputs (i.e. many-to-one, multi-dimensional reaction norms). One alternative is the use of a gene-regulatory network (GRN) models that allow for much more general responses; but this can make analysis difficult. In this work we use a theoretical framework based on logical functions and learning theory to characterize such multi-dimensional reaction norms produced by GRNs. This allows us to reveal a strong and previously unnoticed bias towards the acquisition of simple forms of cell plasticity, which increases their ability to adapt to novel environments. Recognising this bias helps us to understand when the evolution of cell plasticity will increase the ability of plastic cells to adapt to novel environments, to respond appropriately to complex extracellular cues and to enhance developmental robustness. Since this set of properties are required for the evolution of multicellularity, our approach can also contribute to our understanding of this evolutionary transition.

Introduction

Organisms must sense and respond to their environment to develop, survive, and reproduce. Thus, understanding how organisms sense and respond to their surroundings has been a major subject in evolutionary biology from Darwin's times [1–3]. However, during most of the 20th century a simple and convenient schema in which phenotypes are environmentally insensitive (solely determined by genes) was often adopted in the study of evolution [3,4]. Our knowledge of how the sensitivity of the phenotypes to the environment emerges from the developmental dynamics, usually known as phenotypic plasticity, remains far from complete even in the simplest biologically relevant case: the living cell. Since all nucleated cells within a multicellular organism are genetically identical, phenotypic plasticity at the cell level (aka cell plasticity) is necessary for the process of cell differentiation, which in turn is crucial to build up a complex organism composed of many cell types [5,6]. In addition, cell plasticity is also involved in the process of cell de-differentiation, a crucial event for wound healing and regeneration [7].

In cell plasticity, the presence and intensity of different environmental factors (pH, metabolites, morphogens. . .) determines which of a number of potential cell phenotypes will be realised [2]. From a theoretical perspective, the deterministic relationship between environmental factors (inputs) and these specific cell phenotypic states (outputs) observed in a plastic response can be described as a one-to-one mathematical function [8]. Such a one-to-one function (often represented as a one-dimensional reaction norm) suffices to describe the phenotypic plasticity found in most biological systems where the (univariate) phenotype and the (univariate) environmental variable have a simple monotonic or linear relationship (Fig 1A) [8,9].

While one-dimensional reaction norms are a useful heuristic, they cannot accurately describe some important phenomena where phenotypic states are determined by more than

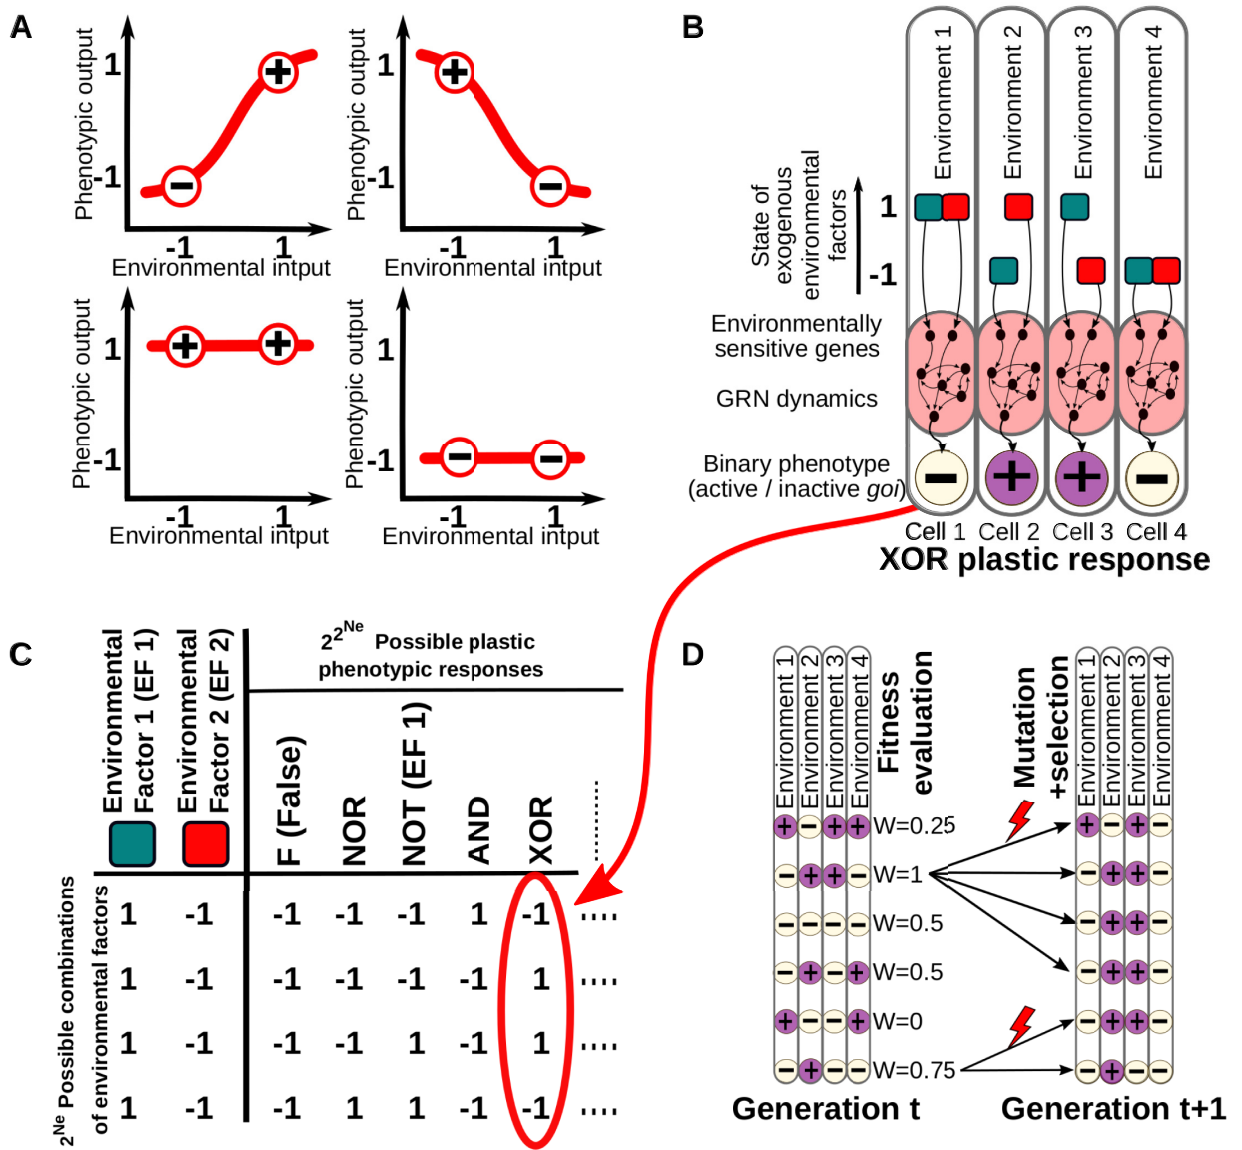


Fig 1. Conceptual depiction of the model. A. Models of phenotypic and cell plasticity often depict a 1-dimensional reaction norm (red line), for a single continuous environmental cue. B. In contrast, we consider different combinations of N_e discrete environmental factors (in this illustrative example, $N_e = 2$, red and blue in the figure) determining a binary cell state. In this case, a number of possible environment-phenotype interactions can be described by means of logical functions derived from Boolean algebra (right). C. Five examples of multi-dimensional reaction norms in which the expression level of some gene is affected by the N_e environmental factors. One of these multi-dimensional reaction norms (e.g. the one described by the “XOR” logical function depicted in B) is set as a target of a continuous GRN-based model. D. The degree of matching between the plastic response of the cells and the expected target multi-dimensional reaction norm determines the fitness (see [Methods](#)). Plastic cells reproduce into the next generations with a probability proportional to their fitness in a mutation-selection scenario.

<https://doi.org/10.1371/journal.pcbi.1006811.g001>

one environmental input [1,2,10]. For instance, differentiating cells use a large number of signalling cascades and other mechanisms (e.g. mechano-transduction, internal biochemical clocks, etc.) as environmental inputs to gather information about when, where, and how to differentiate [7,11]. At the cell level, the output can be described either as the expression profile of many genes or as the expression level of just a single (master) gene that controls the transcription of other downstream genes required for the cell type [10].

In order to describe these many-to-one forms of phenotypic plasticity, we need to use multi-dimensional reaction norms: mathematical functions that generate an output (cell type), given a certain number of input parameters (environmental factors like morphogens) (Fig 1C). The complete set of possible functions is familiar in the field of computer science and described by Boolean algebra (e.g. XOR, AND, NOR, etc.) [12]. Although for the sake of simplicity we have limited the number of inputs to be lower or equal to the number of genes in the plastic cells, the number of possible logical functions escalates exponentially with the number of inputs, providing an ample search space for our experiments (see below). These functions also allow us to describe particular relationships between combinations of environmental cues (inputs) and the phenotypic responses (outputs): any specific multi-dimensional reaction norm of binary inputs and outputs is unambiguously associated with a unique logical function (Fig 1B). Furthermore, this abstraction naturally introduces a complexity measure for multidimensional reaction norms which, as we will show, offers new insights in the study of cell plasticity. Introducing logical functions is also convenient because it allows us to apply principles derived from computer science and learning theory to address some important questions regarding the evolution of phenotypic plasticity. Some previous works have explored how evolving systems can perform Boolean-based input-output relationships where the “calculation” is achieved through a predetermined set of “instructions” [13], through an artificial network with signed states [14,15] or through Boolean GRNs with a fixed, ad-hoc topology [16]. In this paper, the input “calculation” is achieved exclusively through a biological-based and continuous GRN dynamics, which allows us to address, for example: How many signals can a GRN process and how are these signals combined to generate a meaningful phenotypic response? Are some types of cell plasticity more frequent than others? Are all types of cell plasticity functions evolvable? Are some types of cell plasticity more easily evolvable than others? Ultimately these questions are concerned with the structure of the environment-phenotype-maps (EPMs) of cells [17–22]. The first aim of this work is to shed light on these questions and the nature of EPMs by characterizing cell plasticity through this conceptual approach based on logical functions.

In addition, our conceptual approach stresses the similarity between the evolution of cell plasticity and a simple categorisation learning process: cell differentiation requires cells to classify different combinations of environmental inputs (e.g. combinations of morphogens) into a small set of possible categories (cell states). Similar categorisation experiments have been extensively studied in the context of neural networks (NN) [23,24]. These experiments have shown that, during the so called “training phase”, NNs exposed to a number of input-outputs can establish the logical rules that underlie the categorization process, being the rules stored in the NN circuitry [12,25]. Then, in a subsequent “test phase”, NNs apply the same rules to novel, previously unseen inputs to generate an appropriate output. This generalisation ability of a system -i.e., its ability to gather information from previous experience and to use it to offer the appropriate response (e.g. phenotypic, behavioural, etc.) to a previously unseen challenge—is the hallmark of learning [14,26].

The relationship between learning and evolution is more than a casual analogy [12,14,15, 24,26,27]. Although they are not *literally* the same mechanism (e.g., the connection strengths of artificial NNs are deterministically modified during learning [24], whereas variation in the GRN connection strengths occur in an undirected manner through random mutations), learning and evolution can be functionally equivalent at an algorithmic level of abstraction [26,27–29]. That means that, regardless of whether the underlying implementation is deterministic or stochastic, both learning and evolution are exploratory mechanisms that follow local gradients in a space of possible solutions for a problem (the ‘solution space’). Furthermore, both artificial NNs and evolving biological systems feature a non-linear mapping between the ‘model space’

and the ‘solution space’ as is essential for generalisation [26]. In the evolutionary context, this is the mapping between genotype space and phenotype space. By modifying the parameters controlling that mapping during the training phase, the evolution of the developmental process can rearrange which phenotypes are neighbours under genetic mutations and thus can change how phenotype evolution responds to novel selective pressures [14,15,26,29] (or, as here, how evolved plastic phenotypes respond to novel environmental cues).

Although it has been shown that learning principles operate across a plethora of biological phenomena [14,15,26–30], the implications for plastic and differentiating cells performing categorisation learning tasks have not been studied. To ensure the use of appropriate underlying mechanisms, the models we employ in this paper are fully Darwinian—utilising random variation and selection operating on the parameters of a GRN. This type of model has been used many times in previous research (generally without recognising or studying their learning capabilities [16,19]). By recognising, utilising and studying their equivalence to learning systems, we can address questions that have not previously been studied. For instance, we can study the type of generalisation that plastic cells can exhibit and how this depends on the details of the genotype-phenotype map and on the structure of the underlying gene network (GRN).

Previous works have used simplifying assumptions that are common in models of artificial neural networks but not appropriate for natural gene networks [14,15,21]. In this paper, our models make more realistic assumptions regarding the properties of gene-networks, such as non-negative gene concentrations. We also incorporate environmentally sensitive elements in these networks, which are necessary to represent the evolution of plasticity. We then study how natural selection endows individual plastic cells with information about specific environment-phenotype relationships when evolving in heterogeneous environments, which are known to be commonplace in most ecosystems [14,31]. In general, however, cells might not experience *every* possible combination of environmental variables, or may be exposed to a stochastic distribution of such conditions, so that some parts of the multi-dimensional reaction norm remain hidden—not exposed—to selection [16]. We reveal that in these cases plastic cells can sometimes represent the whole reaction norm from incomplete information in the same way that learning systems can induce a complete function from partial data (generalisation) [15].

Generating an adaptive phenotypic response for an environment which has never been experienced in their evolutionary past represents a significant departure from the conventional ‘myopic’ view of natural selection, but is easily interpreted in the light of learning theory [14,26]. We show that such an ability can support phenotypic buffering against noisy environmental conditions, enabling reliable cell differentiation in a context of unreliable morphogen or environmental cues—a phenomenon that may have been important in the evolutionary transition from single-celled free-living organisms to complex multicellular organisms [5,32,33]. Overall, our work suggests that the similarities between learning neural networks and evolving gene networks are not restricted to their structural or functional similarities, but extends to higher level properties such as learning capabilities of the type we will describe. We support these conclusions using biologically plausible models [14,15,34], which suggests that such learning principles are more relevant for real biology than previously demonstrated.

Results and discussion

The core model

The *in silico* experiments presented in this paper have been performed using a GRN-based developmental model with non-symmetric continuous (non-Boolean) unsigned state variables

and explicit environmental factors. The model consists of Ng genes or transcription factors that regulates expression of other genes. The state of the system at a given time is determined by the G vector, containing all the concentration profiles of all the Ng genes and gene products. Genes and gene products have continuous concentrations $g_i \geq 0$, are degraded with a decay term $\mu = 0.1$ and interact with other genes or gene products by binding to cis-regulatory sequences on gene promoters, constituting a gene regulatory network (GRN). The regulatory interactions of this GRN are encoded in a matrix, B , of size $Ng \times Ng$. In B , each interaction strength B_{ij} represent the effect of gene j in the transcription of gene i , acting either as an inhibitor (when $-1 \leq B_{ij} < 0$) or as an enhancer (when $0 < B_{ij} \leq 1$).

Environmental factors are implemented as exogenous cues affecting the levels of gene expression by activating or repressing them. Each i of the Ne environmental factors ($Ne \leq Ng$) affects a single gene with an intensity of $EF_i = -1$ or $EF_i = +1$ (see Eq 2). Although environmental inputs are typically non-binary, a phenotypic transition between two discrete phenotypic states is typically triggered when the environmental cue exceeds a threshold. Our model is restricted to this general case, and assumes that the factor is absent below that threshold and present above it. Environmental factors are not degraded over developmental or evolutionary time (Fig 1). The GRN generates a suitable phenotype by integrating information about both the genetic state and the environment by an iterative process over a number, $devt$, of developmental time steps.

The gene-gene interactions within the GRN follow a non-linear, saturating Michaelis-Menten dynamics, which is a type of Hill function widely used for modelling biochemical binding processes [35]. Other classes of non-linear, monotonically increasing functions have been explored in previous works giving consistent results [36]. Developmental dynamics are attained by changes in gene concentration over developmental time. Thus, the concentration of the gene i in the cell (g_i) changes over time obeying the following differential equation:

$$\frac{\partial g_i}{\partial t} = \frac{R(h_i)}{K_M + R(h_i)} - \mu g_i \quad \text{for} \quad \forall i \in [1, 2^{Ne}] \quad \text{and} \quad \forall j \in [1, N_g] \quad (1)$$

Where $h_i = \sum_{j=1}^{N_g} B_{ij} g_j + EF_j$, given that $EF_j \neq 0 \Leftrightarrow j \leq Ne$ (2)

In the first term of the Eq (1), $R(h_i)$ is the Ramp function ($R(x) = x, \forall x \geq 0$ and $R(x) = 0$ otherwise), which prevents negative concentrations in gene products resulting from inhibiting interactions, and K_M is the Michaelis-Menten coefficient. Without loss of generality, we set $K_M = 1$ (other choices are known to not affect results, [36,37]). The second term of the Eq (1) describes the degradation process that affects every gene in the system. Eq (1) was numerically integrated by means of the Euler method ($\delta_t = 10^{-2}$) over a maximum developmental time of $devt = 10^6$ iterations (arbitrary time units) or until a steady state is reached. This happens when all the normalized gene concentrations remain the same within a threshold of 0.01 over an interval of 1000 developmental time units (i.e., when $|G_{devt}/\max(g(i)_{devt}) - \{G_{devt+1000}/\max(g(i)_{devt+1000})\}| \leq 0.01$). Manually directed simulations showed that when the system does not reach a steady state within this maximum developmental time, it is because it is undergoing either cyclic or chaotic behaviour that does not stabilise. If the GRN does not satisfy this stability criterion, it is assumed to be lethal and therefore replaced by creating a random GRN from scratch or, in the evolutionary simulations (see below), by discarding the last mutation to the GRN and introducing a new one. This procedure limits us to cases where there is viability selection against individuals with unstable networks.

Although the Eq (1) has no explicit noise term, some noise is introduced in the initial conditions in order to break the initial symmetry of the system. Specifically, in $devt = 0$, all gene

concentrations are set to $g_i = 0.1 + \xi$, being ξ a small number randomly drawn from a uniform distribution $\xi \sim U(-10^{-2}, 10^{-2})$. Previous works suggest that small modifications in the equations (e.g. the inclusion of a noise term, the choice of a specific Hill-function, etc.) do not substantially alter the dynamics of these models, thus implying that the developmental dynamics actually relies in the GRN topology and connection strengths [36,37].

The single-trait phenotype is recorded as the gene expression of a gene of interest (g_{oi}), which is different from the environmentally sensitive genes affected by the N_e environmental factors. In many biological examples, specific gene expression levels are known to trigger biochemical, developmental or physiological responses in a binary manner. Above a certain threshold in gene concentration (when the gene is active), the response is triggered, and below the threshold (inactive gene), it is not. Motivated by this, the phenotype of a cell in our model is conceptualized as the binary expression profile of the gene of interest (g_{oi}): $g'_{g_{oi}} = 1 \leftrightarrow g_{g_{oi}} \geq 10^{-2}$ ($= \delta_i$) and $g'_{g_{oi}} = -1 \leftrightarrow g_{g_{oi}} \leq -10^{-2}$. We do not use a lower threshold of zero because some “inactive” gene may have very low but non-zero $g_{g_{oi}}$ values due to the noise introduced in the initial conditions. In this way we treat the phenotypic outcome as a Boolean output. The ordered set of all the $g_{g_{oi}}$ values in all possible environments is denoted by the phenotypic vector P , of size 2^{N_e} . (Note that the activation levels are modelled as continuous values during development but simplified to binary ‘active’/‘inactive’ values at the end of development when values are stabilised).

Although the model is highly simplified and it does not explicitly include space, it captures the most typical mechanisms by which genes regulate one another. Similar models have yielded valuable insights on the evolution of plasticity [38,39], the evolution of modularity [40], developmental memory and associative learning [14,15], the emergence of rugged adaptive landscapes [41], the evolution of biological hierarchy [42], or the role of noise in cell differentiation [43] (Fig 1).

Measuring the complexity of a multi-dimensional reaction norm

In our approach, an environment is defined by a specific combination of (binary) environmental factors (a number N_e of environmental factors can be combined into 2^{N_e} different environments). Since for each environment the cell state can be either $g_{oi} = 1$ or $g_{oi} = -1$, it follows that there are $2^{2^{N_e}}$ potential plasticity functions for N_e environmental factors. The exact relationship between each combination of environmental factors and the binary output constitutes a multi-dimensional reaction norm that is uniquely associated with a logical function (a mathematical input-output map that can be summarised in a so-called *truth table* of length 2^{N_e} , which contains the ordered combinations of all the environmental inputs and the phenotypic (output) vector P). Contrary to approaches based on simple reaction-norms, in which all forms of plasticity are intrinsically equivalent in terms of complexity, the use of logical functions allows us to distinguish between simple and complex forms of plasticity. This is represented by low and high Ω , as defined below ($0 \leq \Omega \leq 1$). As such, Ω is not determined by the number of environmental inputs, but by how complicated is the logic by which they are linked to the cell state [12,44].

In this paper, we use a simple, operational, quantitative and easily computable measurement of the complexity Ω . The function Ω , which is based on other classical measures of information content (e.g. partial information decomposition, [45]), gives us information about the maximal amount of information about the output that is contained in the inputs when these inputs are considered individually. The lesser the amount of information provided by the inputs individually, the greater the complexity Ω , and the larger the amount of computational, digital-like manipulations the system has to apply to the data inputs to obtain the outputs [46,47]. More precisely, the complexity Ω of a given Boolean function by this measure is the

amount of information in the output distribution that is not contained even in the most predictive single-input distribution. In formal terms:

$$\Omega = 1 - \frac{1}{H(P)} \max_{i \in \{1, Ne\}} I(P, EF_i) \quad (3)$$

Where $I(P, EF_i)$ is the mutual information between the distribution of output phenotypes P and an environmental input EF_i , that is:

$$I(P, EF_i) = H(P) + H(EF_i) - H(P, EF_i) \quad \text{for } i \in \{1, \dots, Ne\} \quad (4)$$

In Eqs (3) and (4), $H(X)$ and $H(X, Y)$ denote the (Shannon) informational entropy of the variable X and the joint entropy of the variables X and Y , resp. [48]:

$$H(X) = -\sum_i p(x_i) \log_2 p(x_i) \quad \text{and} \quad H(X, Y) = -\sum_i \sum_j p(x_i, y_j) \log_2 p(x_i, y_j) \quad (5)$$

The left term of the Eq (3) serves to normalise the mutual information with respect to the total information of the output, i.e., the entropy of the output. Because of this normalisation, Ω is not determined for functions with vanishing output entropy (i.e. constant, non-plastic functions; F: always false, T: always true). For non-constant functions (i.e. plastic functions having $goi = 1$ in some environments and $goi = 0$ in others), the function Ω captures the intuition that the simplest functions are the identity functions where the output has the same value of one of the inputs ($\Omega = 0$ for $goi = EF_1$ or $goi = EF_2$), and is maximal when the phenotype is both plastic and very different from all the inputs ($\Omega = 1$ for the non-linearly decomposable functions XOR and XNOR) [12].

Experiment 1a: Natural selection is able to find complex forms of cell plasticity

In a first experiment, we explore whether plastic cells are able to acquire multi-dimensional reaction norms of high complexity. We define cells as able to acquire a specific reaction norm if, when selected to do that (see [Methods](#)), they are able to achieve either the optimum fitness ($W = 1$) or a relatively high performance ($W > 0.9$). Based on the recorded fitness over evolutionary time (see [Methods](#)), our first experiment shows that cells are able to acquire multi-dimensional reaction norms of high complexity (the highest possible complexity $\Omega = 1$ for the Ne considered, [Fig 2](#)), although the data also suggest that there are strong dependencies on the three parameters considered.

For instance, the amount of time required to adapt increases rapidly as the environmental dimensionality increases (See upper panels in [Fig 2](#)). Furthermore, there exists a minimal GRN size required to generate a given multi-dimensional reaction norm. This is likely because, contrary to some neural network (NN) modelling approaches, GRNs have strictly non-negative states (a gene concentration cannot be negative) which prevents a straightforward digital-like computation of the inputs (see [Discussion](#)). The results also suggest that this minimal size depends on the complexity of the function required; the more complex the function the larger the network required to achieve it. All else being equal, plastic cells take longer to evolve complex multi-dimensional reaction norms than simpler ones.

This initial experiment shows that, given enough time and suitable selective pressures, biologically plausible GRNs above a certain size are capable of producing complex relationship between environmental inputs and phenotypic states, including those with maximal $\Omega = 1$ ([Fig 2D](#)). How the evolutionary time required to evolve complex functions with $Ne > 4$ remains to be investigated ([Fig 2B](#)).

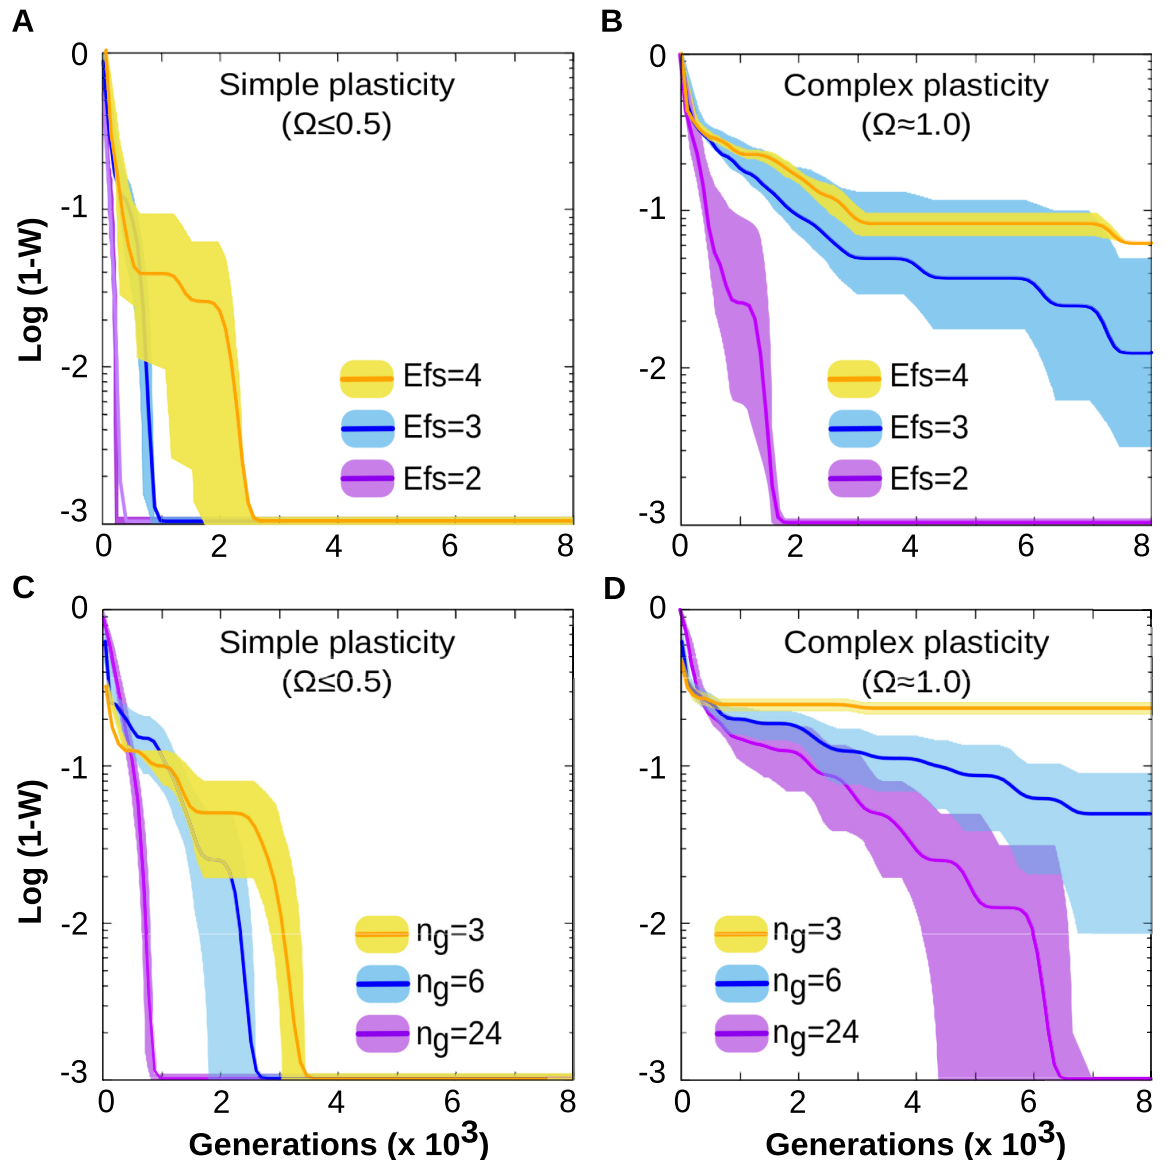


Fig 2. Plastic cells are able to evolve forms of plasticity of arbitrary complexity. (A and C) Left panels: Different environmental inputs can be easily integrated in a phenotypic response by means of a simple forms of cell plasticity (e.g. AND, NOT functions). Evolving cells can also evolve complex forms of complexity (e.g. XOR, XNOR functions), but they take longer to do this. In general, evolving a specific form of cell plasticity is faster when the number of environmental factors involved is low (Upper panels A-B), and/or when the size of the GRN is large (Lower panels C-D). Below certain thresholds in the number of environmental factors and in the GRN size, (which are complexity-dependent), the system fails to accurately represent the function (see yellow lines in the right panels). Each line shows the average and standard deviation over 30 replicates. Each replicate has the same initial conditions (an “empty network” with $B_{ij} = 0, \forall i, j \in [1, Ng]$) but traces a different evolutionary trajectory (different mutations). For each replicate, the target function is randomly drawn from the set of possible functions with the same N_e and Ω . $N_g = 12$ in A-B and $N_e = 3$ in C-D.

<https://doi.org/10.1371/journal.pcbi.1006811.g002>

Experiment 1b: An alternative mechanistic implementation of cell plasticity enhances evolvability

There are at least two mechanistic ways in which an environmental signal can affect GRN dynamics: by affecting the gene expression (either by activating or repressing it), or by modifying the strength of regulatory interactions. In this work, we have explored these two

possibilities, which we refer to as the classical and the second-order implementations, respectively. Recent works suggest that biochemical mechanisms consistent with the second-order implementation are common in nature and are involved in biological phenomena like proteins with intrinsically disordered domains [49,50] or non-deterministic GRNs [51]. Other examples include temperature, which causes huge phenotypic effects by affecting the morphogen diffusion or the ligand-receptor kinematics, but without altering the cellular concentration of these elements. It is worth noting that factors affecting the strength of regulatory interactions coexist with factors that affect regulatory mechanics (e.g. pheromone-like chemicals).

These GRNs whose topology is not fixed but dynamic over developmental time have been implemented by transforming the B_{ij} matrix into a second-order matrix (tensor) B_{ijk} . The extra dimension “ k ” determines which genes contribute to the establishment of the interaction strength between the gene products i and j . Therefore, the h_i term of the Eq (2) is now described as:

$$h_i = \sum_{j=1}^{N_g} g_j \sum_{k=1}^{N_g} B_{ijk} g_k + EF_k, \quad \text{given that } EF_k \neq 0 \Leftrightarrow k \leq Ne \quad (6)$$

In this second-order implementation, each environmental factor also contributes to dynamically modify the strength of regulatory interactions with an intensity of $EF_k = -1$ or $EF_k = +1$.

In order to assess if the mechanistic manner in which the environment informs development can affect the evolution of cell plasticity, we reproduce the results of experiment 1a but by using instead the second-order implementation (except for this one, all experiments in this work were carried out using the classical implementation). As Fig 3 indicates, the results from Experiment 1a hold qualitatively for both classic and second-order implementations. However, cells equipped with GRNs capable of second-order dynamics can evolve such plasticity much faster than cells with classic GRNs. Results also show that the greater differences in the adaptive rate between the classic and second-order implementations occur for complex forms of plasticity (Fig 3B). This can be explained by considering that in complex logical functions the phenotypic effect of one environmental factor EF_1 is dependent on the state of EF_2 (higher level correlations). This dependency is naturally captured by the second-order implementation, where the environmental factors act as modulators of the genetic effects rather than as determinants of genetic states (Fig 3A). In addition, the dependence between the time required for adaptation (i.e., reaching a fitness of $W = 0.95$) and GRN size becomes more linear under the second-order implementation, thus relaxing the necessity of very large GRNs to generate complex form of plasticity (Fig 3B, red lines).

Overall, the results summarised in Fig 3 suggest that when environmental factors affect the strength of gene-gene interactions, rather than the gene concentrations, the ability of biological systems to evolve complex forms of plasticity improves. Notice that this claim does not imply the mechanisms of cell plasticity consistent with the second-order implementation to be more common than those based on classic GRNs. It simply states that, when those mechanisms are available, the ability of cells to evolve complex forms of cell plasticity is greatly improved.

Experiment 2: Simple forms of cell plasticity are far more abundant than complex ones

One of the possible interpretations of the results of the first experiment is that evolving cell plasticity associated with complex functions is more difficult because the number of GRNs performing these functions is low. That is, given any GRN generating a particular multi-dimensional reaction norm, most of its mutational neighbourhood will most likely produce either the same form of plasticity or a simpler one. In order to confirm this hypothesis, we

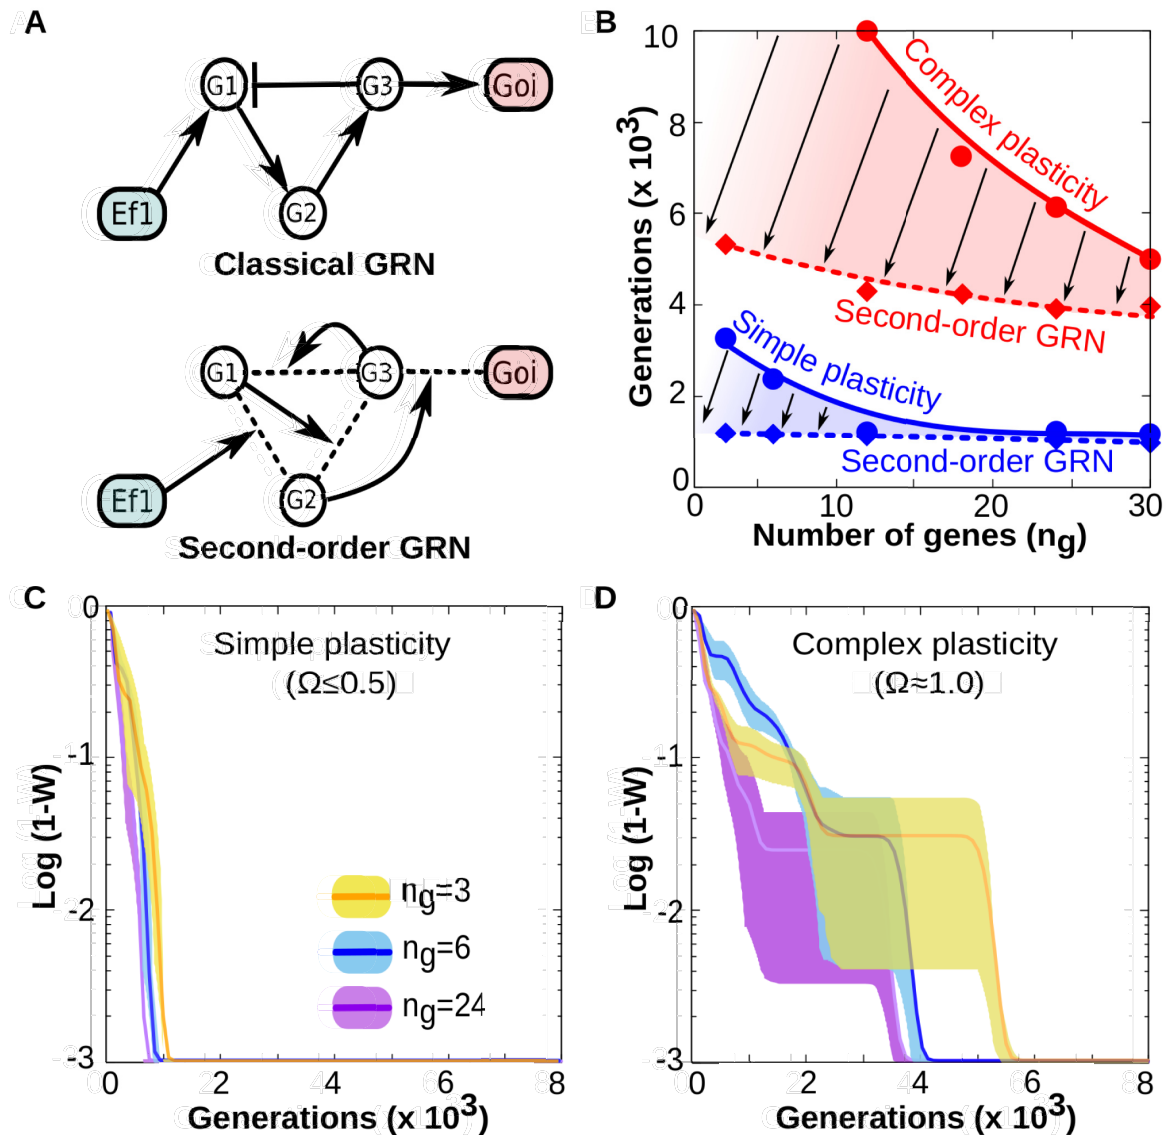


Fig 3. Second-order implementation of cell plasticity enhances evolvability. A) Contrary to “classical” GRNs in which the network topology remains fixed over developmental time, second -order (tensor-based) GRNs have dynamical topology: each gene-gene interaction strength is determined by the concentration of other genes (g_1, g_2 , etc) and environmental factors. In both cases the final phenotype is recorded as the binary state of a gene of interest (goi). B) Comparison between “classical” and “second-order” GRNs in terms of their evolvability (how many generations they need to achieve a fitness $W \geq 0.95$ when they are selected to represent specific reaction norms). The panel shows that everything else being equal (number of genes, function complexity), GRNs with dynamical topologies (dashed smoothed lines) require much less generations to evolve any reaction norm than classical GRNs (solid smoothed lines). Thus, GRNs capable of second-order dynamics exhibit higher evolvability than classic GRNs (being the difference more patent for complex forms of plasticity). C-D) As in the classical implementation, evolving a specific form of cell plasticity is faster when the size of the GRN is large, although this relationship is much more linear (compare with Fig 2C and 2D) $N_e = 3$ and 30 replicates for all panels. Each replicate has the same initial conditions (an “empty network” with $B_{ij} = 0, \forall i, j \in [1, N_g]$) but traces a different evolutionary trajectory (different mutations). For each replicate, the target function is randomly drawn from the set of possible functions with the same N_e and Ω .

<https://doi.org/10.1371/journal.pcbi.1006811.g003>

performed an unbiased scanning of the parameter space to investigate how the different types of plasticity are spread over the theoretical space of all possible GRNs and how abundant each multi-dimensional reaction norm is. To do this, we generated a large number ($\approx 10^5$) of

random GRNs with a size of ($3 < Ng < 24$), so that two of their genes were environmentally sensitive to EF_1 and EF_2 and a third, different gene was set as *goi*. In order to obtain representative networks for different values of the parameters considered (average strength of connection and percentage of non-zero connections, which are proxies for the L_1 and L_2 regularisations parameters respectively, see costly connections below) the values of the B_{ij} matrix for a given GRN were given, with $p = U \sim (0,1)$, non-zero values randomly drawn from a normal distribution $N \sim (U \sim (-1,1), \sigma)$ with $\sigma = 0.1$. That way we avoid the convergence of the average absolute strength (L_1) to a value of 0.5, derived from the central limit theorem.

The results, summarized in Fig 4A, show that indeed most of all the possible GRNs do not exhibit cell plasticity at all. From the subset of GRNs that show plasticity, we observe that complex plastic response functions are one order of magnitude less frequent than simple plastic response functions. Thus, the relative frequency of a given multi-dimensional reaction norm is inversely related to the complexity of the logical function which describes it. The histogram in the Fig 4A shows that the frequency-complexity relationship approximates a logarithmic function. Notice that these different frequencies do not arise from the relative frequency of each type of function in the mathematical space of all possible logical functions (their probability functions are plainly different, see dashed line in Fig 4A), but emerge as a derived property of GRN dynamics.

We also examined the sub-volumes of the parameter space where each type of multi-dimensional reaction norms is more likely to be found. Fig 4B–4D show that, except for very complex reaction norms ($\Omega = 1$, yellow dots in the plots), all are randomly scattered across the parameter space. The asymmetric sorting observed suggests that GRNs accounting for complex multi-dimensional reaction norms must have a minimum number of elements (Fig 4C) and dense connectivity with strong and disparate gene-gene interactions (Fig 4A and Fig 4D). In nature, selective pressures, such as a metabolic cost of keeping densely connected GRNs, may drive the system away from these regions enriched with complex forms of cell plasticity, making them evolutionarily inaccessible even if they are beneficial [15].

Together with the results from experiment 1, this bias towards the use of simple forms of plasticity seems to imply that plastic cells evolve in search spaces of reduced dimensionality: they are more likely insensitive to many of the inputs they are exposed to, and preferentially establish simple correlations between the remaining ones. This scenario of reduced search spaces has been proposed on other theoretical and experimental grounds [10,16,43,52].

Experiment 3: Simple forms of cell plasticity are evolutionary attractors

We next test whether evolutionary changes from simple to complex reaction norms are as likely as changes in the opposite direction. That is, we checked if the transitions between the different forms of cell plasticity are isotropic ($p(a \rightarrow b) = p(b \rightarrow a) \forall \Omega_a > \Omega_b$) or not. In order to do that, we first evolved a number of GRNs ($n = 30$) to produce a specific multi-dimensional reaction norm with an associated complexity of Ω_a (when the GRN reached a $W = 1$, the simulation was stopped). Since this process is analogous to the “training phase” of artificial neural networks, these evolved networks are referred hereafter as “trained networks”. Then we clone the trained network and force each clone to evolve again towards a new form of plasticity, now characterized by a complexity Ω_b . We recorded the inverse of the number of generations required to attain again a $W = 1$ as a proxy of the likelihood of the transition $a \rightarrow b$, which can be interpreted as the system’s evolvability.

The results show that, in general, transitions between different complexity classes are not isotropic (Fig 5A): the number of generations needed to evolve simpler reaction norms is smaller than the number of generations needed to evolve more complex reaction norms (p

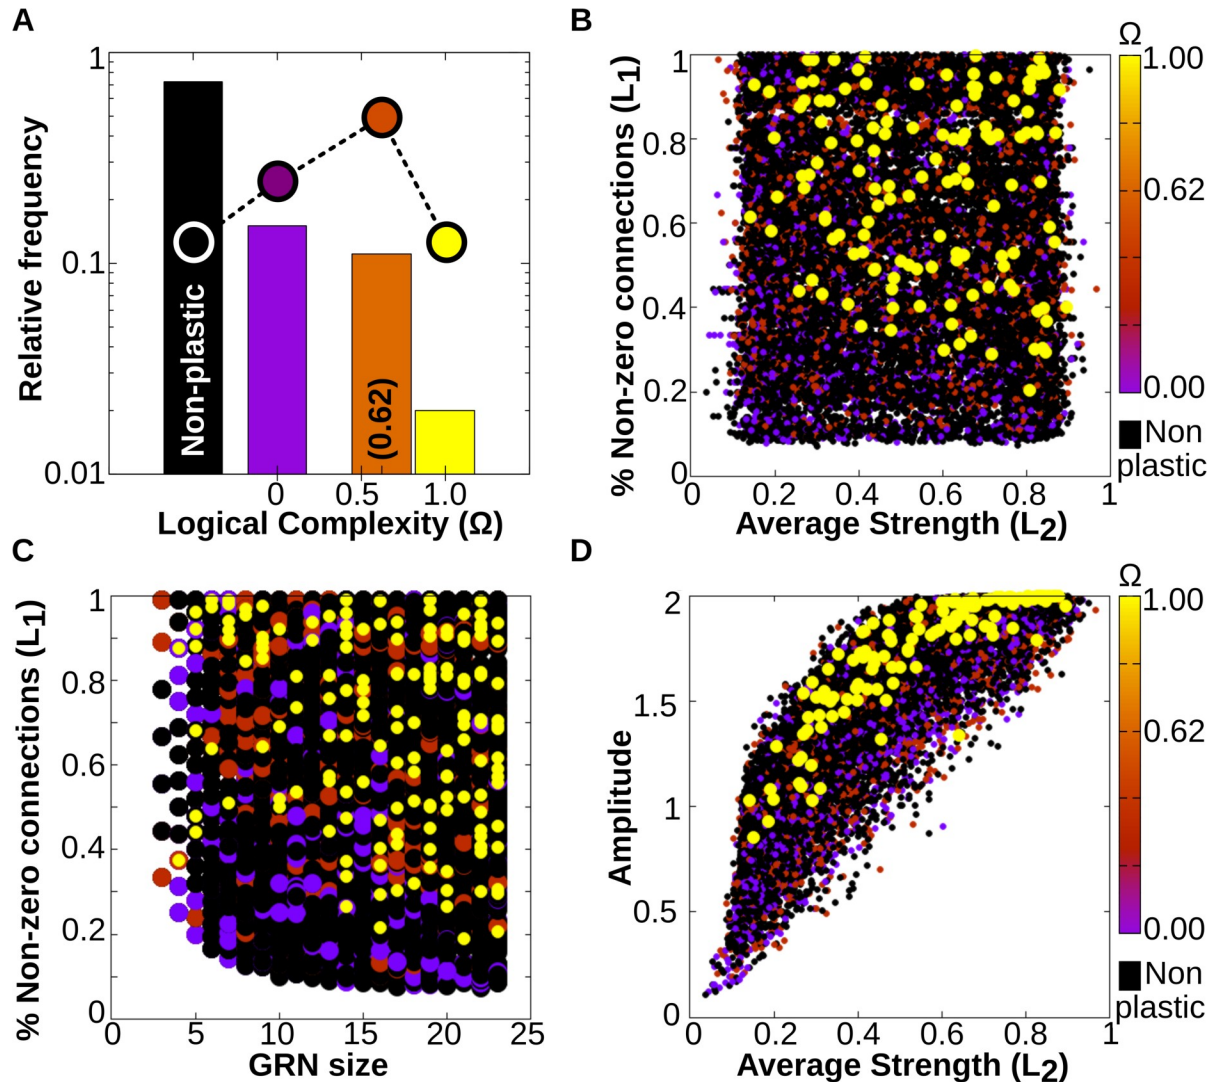


Fig 4. Exploring the morphospace of cell plasticity. A. (Log) Relative frequency of different types of cell plasticity according to their complexity Ω in a vast GRN space. Black column: no plasticity (the phenotypic state is purely determined by genes); purple column: phenotypic state directly determined by just one of the N_e environmental factors; orange column: phenotypic state determined by simple combinations of the N_e environmental factors (linearly decomposable functions) and yellow column: complex forms of cell plasticity associated with non-linearly decomposable functions (XOR, XNOR; see SI). Dashed line and dots represent the relative distribution of each family of logical functions in the mathematical space. We see that although the number of simple and complex functions that exist is approximately equal, GRNs produce simple functions much more often. B. Complex forms of plasticity (yellow dots) arise preferentially in densely connected networks with strong gene-gene interactions. C-D. Complex forms of plasticity require minimum GRN sizes and disparate ($0 << |B_{ij}| < 1$) values in the GRN connections. In B-D panels, the % of non-zero connections and the average strengths of the GRNs correspond, respectively, to the parameters L_1 and L_2 used for regularisation procedures (see Eqs 9 and 10). In general, less complex functions with $\Omega < 1$ are evenly distributed in the parameter space, and are not associated with specific GRN topologies. ($N_e = 2$; $\approx 10^5$ points).

<https://doi.org/10.1371/journal.pcbi.1006811.g004>

$(a \rightarrow b) > p(b \rightarrow a) \forall \Omega_a > \Omega_b$). Thus, although experiment 2 shows that complex forms of cell plasticity can evolve, experiment 3 shows that evolutionary changes to simpler forms of cell plasticity are generally favoured (i.e, they can happen in a smaller number of generations, all else being equal). In general, our data suggest that, given two different multi-dimensional reaction norms a and b , the speed of transitioning between them is proportional to the differences between their associate complexities ($\Omega_a - \Omega_b$). If both reaction norms belong to the same complexity class ($\Omega_a = \Omega_b$), the speed of transitioning between them is inversely proportional to

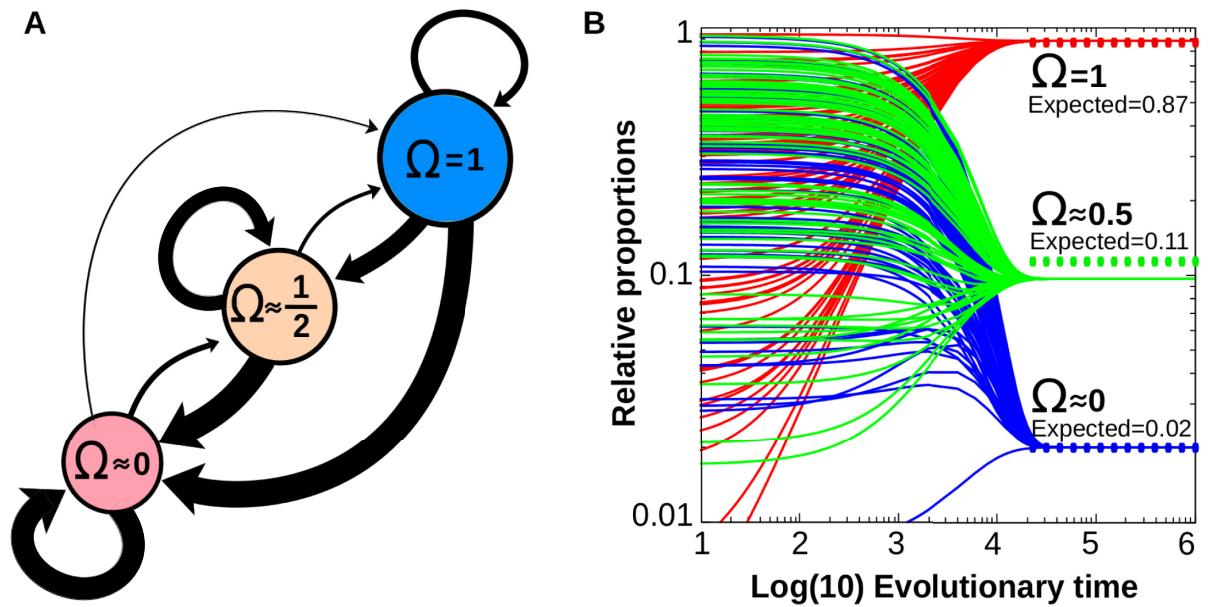


Fig 5. Evolvability assays. A. Transitions from complex forms of cell plasticity to simpler ones are evolutionary favoured. The width of each arrow is inversely proportional to the number of generations required for the transition from a specific multi dimensional reaction norm to another one (the wider the arrow the more probable the evolutionary change). The area of each circle is proportional to the average evolvability of each multi dimensional reaction norm. B. Predicting the relative frequency of each type of cell plasticity from transition probabilities. Numerical simulations for the expected long-term distribution of each type of plasticity derived from the probability transitions shown in A (see Eq (7)). Convergence, steady state values (0.8833, 0.0966 and 0.0201) are similar to those expected (dashed lines) according to the random scanning of the GRN space (Experiment 2), suggesting that the differential evolvability (the different probability transitions between different forms of cell plasticity) emerges from the drastic differences in the relative frequency of each type of cell plasticity in the GRN space.

<https://doi.org/10.1371/journal.pcbi.1006811.g005>

their complexity class (Fig 5A, looped arrows). More formally: $p(a \rightarrow b) \propto \Omega_a - \Omega_b \forall \Omega_a \neq \Omega_b$ and $p(a \rightarrow b) \propto \Omega_a^{-k} \forall \Omega_a = \Omega_b$.

The average number of generations required to go to any other reaction norm (including those belonging to its own complexity class) is marginally increased for complex multi-dimensional reaction norms compared to simple reaction norms (see relative sizes of the circles in Fig 5A). Thus, there appears to be a consistent evolutionary bias towards evolving simpler multi-dimensional reaction norms irrespective of the initial reaction norm considered.

We hypothesize that the differences in transition probabilities between forms of cell plasticity emerges from the relative frequency of each type of cell plasticity in the GRN space (see Experiment 2). We tested this by numerically calculating the expected long-term distribution of each type of plasticity derived from the probability transitions shown in Fig 5A. These probability transitions were introduced in the differential Eq (7):

$$\frac{\partial \Omega_i}{\partial t} = \sum_{j=1}^{\Omega_{classes}} p(\Omega_j \rightarrow \Omega_i) \cdot \Omega_j \tag{7}$$

Where Ω_i is the relative frequency of the type of complexity i (amongst the $\Omega_{classes}$) and $p(\Omega_j \rightarrow \Omega_i)$ is the transition probability between the Ω_j and the Ω_i classes. Eq (7) was iterated until steady-state values in $t \rightarrow \infty$, which were then compared with those resulting from the relative frequencies of experiment 2. The calculations show that, when plastic cells change from one multi-dimensional reaction norm to another according to the probability of that transition (found in Experiment 3), the relative frequency of each type of cell plasticity converges to a steady-state, equilibrium value over long evolutionary timescales (Fig 5B). These values are

almost identical to those yielded by the random scanning of the GRN space (Experiment 2), suggesting a causal connection between the probability transitions and the relative abundance of each type of cell plasticity (notice that these values come from different experimental setups, so they could be different).

Together, experiments 1 to 3 demonstrate a strong and previously unappreciated bias towards establishing the simplest form of cell plasticity. These simple forms of plasticity are exemplified by phenotypic responses which are triggered by simple (e.g. linear, additive) combinations of the environmental cues.

Experiment 4a: Plastic cells perform adaptive generalisation of simple plastic responses

In this experiment, we test whether cells exhibiting different classes of cell plasticity are able to use their past evolutionary experience to better evolve to a new environmental challenge (that is, if they exhibit some learning capabilities [26]). We recreate classical categorisation experiments, widely used in learning theory [12,15], by which the system has to exploit regularities observed in past situations to offer appropriate responses in novel cases (generalisation). For plastic cells, generalisation means that evolution is able to infer the whole multi-dimensional reaction by being exposed to a fraction of it (a subset of the input-output relationships). We know from experiment 1 that plastic cells can learn a complex reaction norm when they are exposed to all of its points, so we now check if evolution can also learn and reproduce a target reaction norm when there is missing information. In the approach we follow here, a given multidimensional reaction norm contains 2^{N_e} points, which are summarised in a truth table of the same length (Fig 1A). To model a scenario with missing information, we set a training phase in which evolving cells can only sense a random fraction of the complete truth table. The “training set size” (TS) denotes the number of environments (rows of the truth table) that are available for cells during the training phase ($TS < 2^{N_e}$).

As Fig 6A–6C shows, the ability of evolution to generalize depends on the complexity of the target multi-dimensional reaction norm. When natural selection experiences only a fraction of all possible environments, it results in plasticity that interprets the environmental correlations in the simplest possible manner. That is, when some information is missing, natural selection finds the simplest forms of plasticity compatible with the input-output relationships that it has experienced during the training phase. This finding is in agreement with the bias shown in our experiments 1–3 (see points with $\Omega < 0.5$ in the left side of Fig 6C). This bias towards simplicity is advantageous when the target reaction norm has low complexity (Fig 6A): since plastic cells generalise a simple reaction norm in unseen environments, they have greater chances to fit the adaptive (target) function. Thus, under some circumstances (low complexity functions), the inductive bias exhibited by plastic cells allows them to perform better than using random responses in novel environments (Fig 6A and S1 Fig, blue line above the green line). By random responses we mean that for every point of the reaction norm for which the system has no previous information, the phenotypic response has equal probability of being +1 or -1.

Conversely, when the structure of the problem requires fitting complex ($\Omega > 0.5$) multi-dimensional reaction norms, cells cannot accurately predict the phenotypic response required in the new environments (blue line in Fig 6A and 6B).

In both simple and complex cases, the differences between the performed plastic response and the random response are larger for $TS = 2^{N_e}/2$, that is when the system can access just half of the available information. Obviously, also in both cases generalisation capability of evolution improves when the size of the training set increases, thus providing more information to natural selection, up to the limit case in which $TS = 2^{N_e}$, which was the scenario explored in experiment 1.

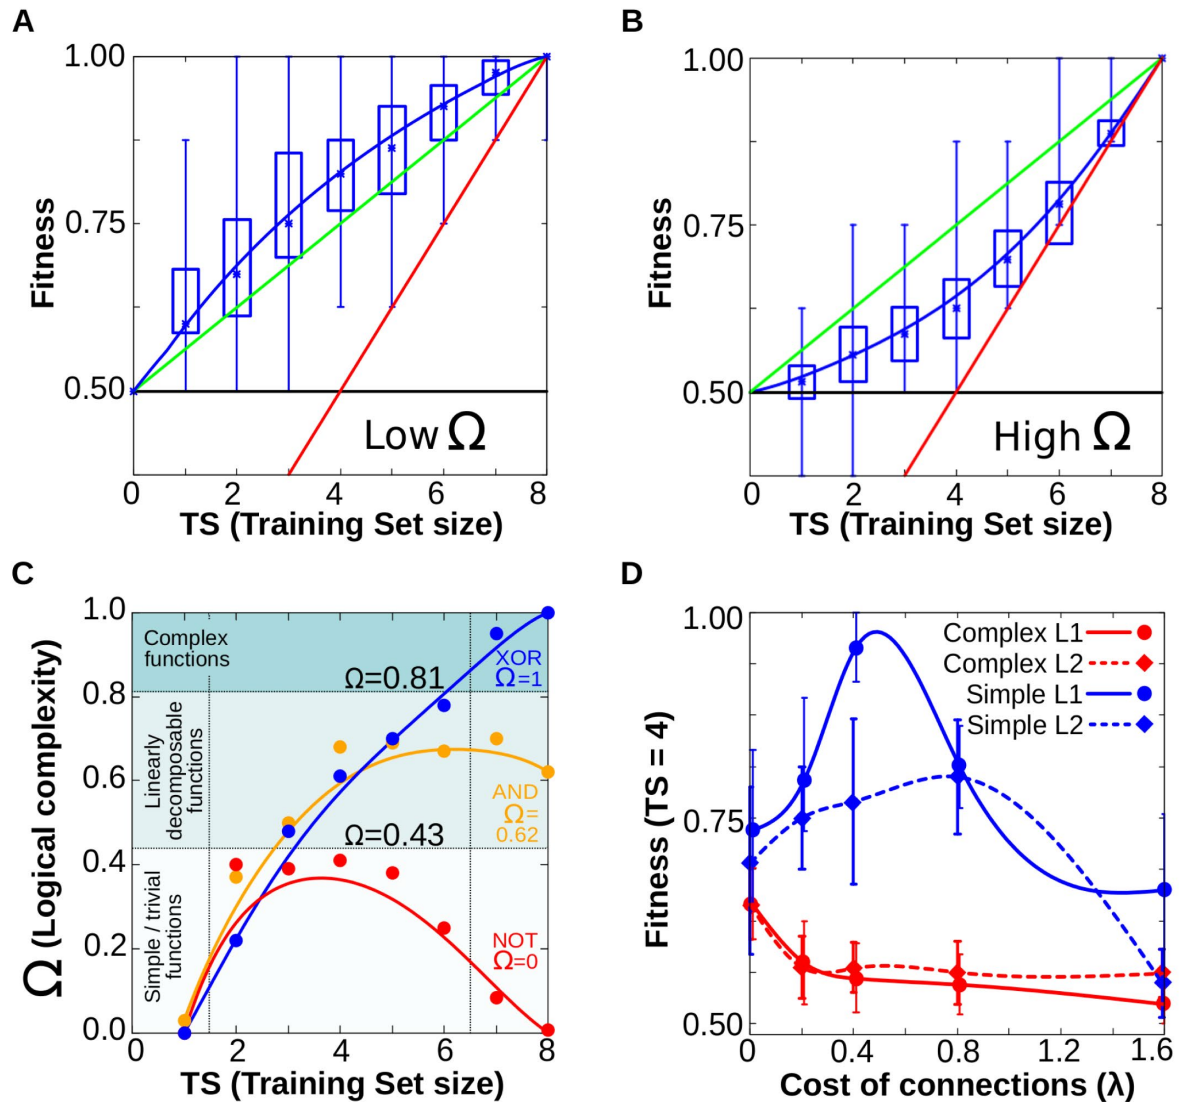


Fig 6. Generalisation experiments. A. When cells evolve simple forms of plasticity, they are able to generalise, performing better than chance in previously unseen environments (red line represents the information provided ($x = y$), green line represents the expected performance at random ($2^{N_e} - TS$)/2; and blue line the degree of matching between the resulting phenotypes and the expected ones for a given function). Notice that blue line runs consistently higher than the green line (random response). B. Similar experiments yield poor performances when cells have to evolve complex forms of plasticity (blue line below the green one, see main text). C. When cells have incomplete information about all possible environments (left), they acquire preferentially simple forms of cell plasticity. Acquiring complex forms of cell plasticity beyond linearly-decomposable functions (see SI) requires full information (blue line). In this panel, the complexity Ω of non-plastic functions was considered to be zero. D. Generalisation experiment under L_1 and L_2 regularization. For this experiment, plastic cells are only exposed to half of the possible environments ($TS = 4$), that is the point where their inductive bias makes the result most different from randomness. We record how the ability of cells to generalise (to approach the maximum fitness of $W = 8$) changes as the cost of connections for the GRN increases (see Eq (8)). We do this for L_1 and L_2 regularization (solid and dashed lines respectively) which favour sparse and weak connections respectively. We also show data for both complex ($\Omega \approx 1$, red lines) and simple ($\Omega < 0.5$, blue lines) plastic responses. The plot shows clearly how including a cost of connections improves drastically the performance of cells in generalisation experiments when cells are asked to evolve towards simple forms of plasticity. For all panels $N_e = 3$, $N_g = 12$, $n = 30$ replicates, standard deviation as boxes, min and max values as error bars.

<https://doi.org/10.1371/journal.pcbi.1006811.g006>

Notice that for these experiments, the target plasticity function for each replicate was drawn from a subset of functions, i.e. those exhibiting the desired Ω , not from the whole ensemble of all possible logical functions of N_e inputs. Therefore, the tendency of selection to evolve simple

reaction norms in this experiment cannot be attributed to their large statistical availability, but rather to biases resulting from GRN dynamics.

In summary, when the adaptive phenotypic response is a simple function of the environmental inputs, cells will produce adaptive phenotypes in the new (not previously experienced) environments better than expected by random completion of previously unseen rows of the truth table (Fig 6A, S1 Fig). This feature does not derive from previous selective pressures alone (i.e. it goes beyond rows that have been observed by natural selection in the past), but can be viewed as resulting from an inherent bias of evolving GRNs. That is, from the set of plasticity functions that are compatible with past selection, evolved GRNs will ‘over sample’ simple plasticity functions compared to complex ones. In learning terms this is what is meant by an “inductive bias”, necessary for generalisation [15].

In an additional experiment, we explore whether some selective conditions could enhance generalised phenotypic responses. In the light of Experiment 2, there seems to be a dependency between the complexity of the plastic phenotypic responses and the GRN topology (e.g., complex multi-dimensional reaction norms are produced more often by plastic cells having densely connected GRNs). Thus, we examine an evolutionary scenario in which the generalisation experiments must be performed by GRNs with a particular topology. Specifically, we force GRNs to acquire a certain topology (e.g. a weak connectivity) while evolving in the training phase of a generalisation experiment. We do this by implementing a connection cost in the evolving GRNs, so that the connectivity of the GRN contributes to the cell’s fitness (see Eq (8)). Although connection costs in the GRNs can not typically be demonstrated experimentally, such a cost is biologically meaningful since densely inter-wired GRNs are known to be selected against for a number of reasons, including the intrinsic metabolic cost of synthesizing a larger number of organic molecules, the lack of efficiency in the substrate-ligand interactions between gene products derived from thermodynamic considerations, and because the GRN dynamics is more likely to exhibit chaotic behaviour [15,37,41]. Thus, although there exist some evolutionary advantages in having redundant elements in the GRN (e.g. increased robustness [19]), the selection against unnecessarily complex GRNs is biologically consistent.

In addition, introducing a connection cost (aka parsimony pressure or regularisation pressure [15]) is a widely used procedure in learning theory and artificial intelligence that is known to alleviate the problem of over-fitting [12,14,15], thus providing us an opportunity of testing the formal correspondence between learning principles and cell plasticity. This procedure can be easily implemented by making the GRN connections costly, with a direct effect on fitness:

$$W_t = W_p - \lambda \cdot CC \quad (8)$$

Where W_p is the partial fitness scored as in Eq (11) (see Methods) and λ is the relative weight of the cost of connections (CC) in determining the total fitness W_t . The cost of connections (CC) was implemented in two different ways. The first one (known in computer science as L_1 regularisation) consists basically of limiting the complexity that the network can attain (henceforth favouring sparse connectivity in the GRN). This is easily implemented by applying a direct selective pressure in the magnitude of gene-gene interactions, which effectively decreases the number of regulatory connections in the GRN [15].

$$CC(L_1) = \frac{\sum_{i=1}^{n_g} \sum_{j=1}^{n_g} |W_{ij}|}{n_g^2} \quad (9)$$

The second procedure (aka L_2 regularisation) favours GRNs with weak connections (i.e. regulatory interactions of small magnitude [15]). This is implemented as a selective pressure in

the square of the strength of gene-gene interactions:

$$CC(L_2) = \frac{\sum_{i=1}^{n_g} \sum_{j=1}^{n_g} W_{ij}^2}{n_g^2} \quad (10)$$

As Fig 6D shows, in these experiments L_1 and L_2 procedures have a qualitatively similar effect on the ability to generalise. However, the effect depends strongly on the complexity of the multi-dimensional reaction norm that cells have to evolve (S2 Fig). For complex forms of plasticity (red lines in Fig 6D), introducing a cost of connections makes the generalisation response worse (approaching the minimum value of $TS = 2^{N_e}/2 = 4$, where the system output contains just the information provided during the training stage). This happens because introducing a cost in the GRN connections pushes the system towards regions of the GRN space enriched with weakly and sparsely connected GRNs (where no complex forms of plasticity can arise), so the system cannot fit the complex target (known as “underfitting” [15]).

The opposite is observed when the evolved cells generalise over simple forms of plasticity. In that case (blue lines in Fig 6D), having sparse and weakly connected GRNs increases the likelihood of evolving cells that perform simple forms of plasticity, which increases their performance under previously unseen environmental conditions (i.e., generalisation). However, when the connection cost becomes too high (beyond an optimum λ of 0.5 for L_1 and 0.8 for L_2), the enhanced performance disappears. That is probably related to the fact that even to perform the most basic computations, GRNs require a minimally dense topology: extremely simplified GRNs do not perform any kind of computation at all (S2 Fig). Summarising, our set of additional experiments suggest that costly connections amplify the performance in generalisation experiments. This suggests that some learning principles widely used in the field of artificial NNs have direct functional equivalence in the context of cell plasticity.

Experiment 4b: Learning principles explain improved developmental stability in a multicellular context

In complex multicellular organisms, the phenotypic state of the individual cells is determined by a complex combination of environmental cues and endogenous signals such as morphogen gradients [7]. In order to produce non-trivial and spatially organized developmental patterns, differentiating cells need to respond appropriately to these complex signals. In principle, it could be the case that every possible response (cell-state) has been the explicit target of past natural selection. However, the previous experiment suggests that this is not always necessary. If the target pattern results from simple morphogen combinations and natural selection has occurred over enough morphogen-phenotype combinations, the evolved GRNs will fill-in remaining regions of the multi-dimensional reaction norm better than by chance. To illustrate the significance of this generalisation capacity in a multicellular differentiation context, we set a 2D embryonic field of 1250 cells which can gain positional information from the asymmetric spatial distribution of a few different morphogens (Fig 7). For the sake of simplicity, we consider that the cells involved can receive but not release morphogens, and that the (exogenous) morphogens are spatially distributed along simple gradients. Importantly, none of the individual morphogens contains enough information to produce the final developmental pattern. The pattern emerges by integrating the morphogenetic cues according to a relatively simple composite logical function (Fig 7A). Although natural GRNs co-evolve with the morphogens and the pattern themselves, the latter often evolve at slower rate than the GRNs [53,54]. For the purposes of this illustration, we consider a scenario where these differences in evolutionary rates are sufficiently large so that both the pattern and the gradients are effectively constant on the timescale simulated. All these simplifications enable us to isolate the problem

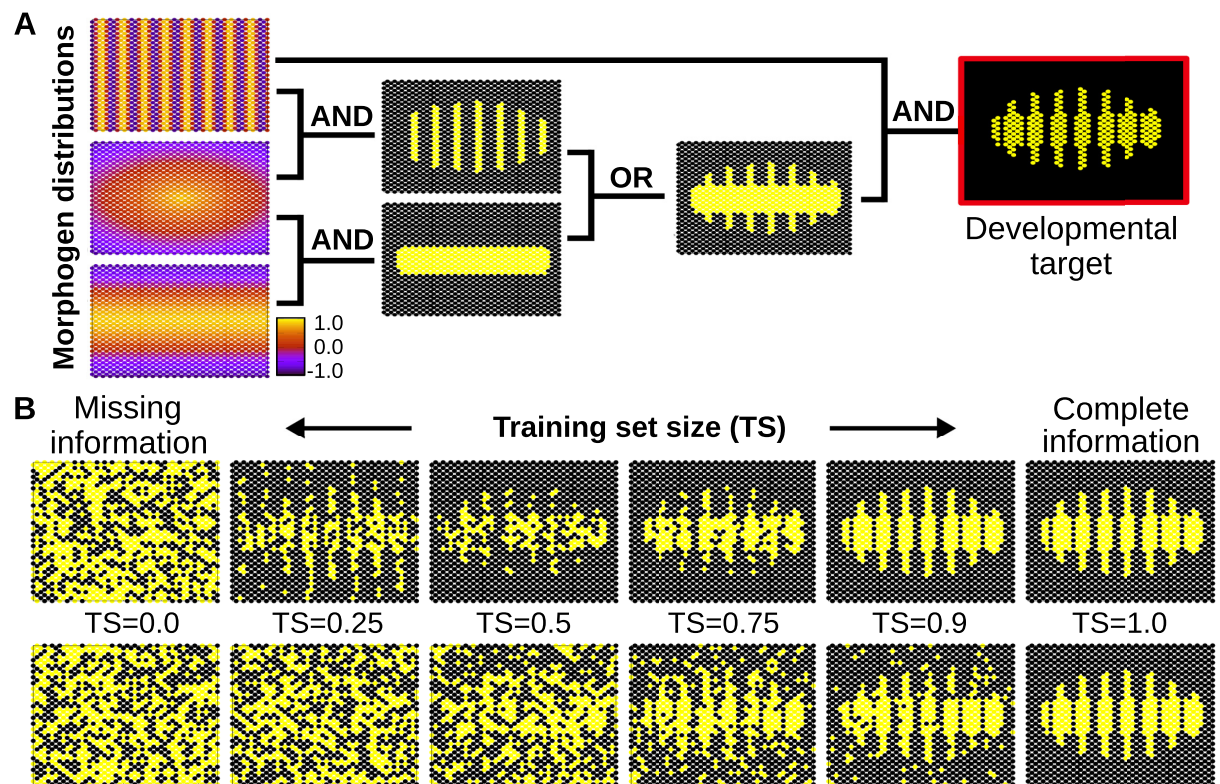


Fig 7. Developmental implications of learning principles in the evolution of cell plasticity. A) Simple spatial distribution of three morphogens over a two-dimensional (50x25) field of cells (similar gradients are commonly found in the early developmental stages of many organisms). In this example, these three environmental signals are integrated by individual cells (according to relatively simple logical functions) to form a segmented developmental pattern. The complete plasticity function is represented by a truth table of 32 rows specifying whether the relevant cellular output should be on or off for each possible combination of morphogens (see [Methods](#)). B) Individual cells are evolved under selective conditions that expose them to some rows of this table but not all. Under this scenario of missing selective information, cells equipped with real GRNs perform by default a simple logical integration of the morphogen cues, which in this example results in a much more robust and uncertainty-proof developmental process (upper row). When the training set contains all the information necessary to completely define the plasticity function, natural selection finds a GRN that calculates this function accurately (right). Of course, when past selection contains no useful information, the phenotype of the GRN is random (left). In between we can see that the generalisation capability of the evolved GRN 'gives back more than selection puts in' - i.e. the phenotype produced (top row middle) is visibly more accurate than the training data experienced in past selection (bottom row middle). This is quantified in [S1 Fig](#). In the bottom row, cells do not exhibit any bias towards simple functions, and therefore they acquire any random form of plasticity compatible with the previously experienced environmental inputs. This randomness in the plastic response fails to fit the required (target) function, preventing generalisation (i.e. the system needs to evolve in a full-information scenario to produce the required phenotype, see main text).

<https://doi.org/10.1371/journal.pcbi.1006811.g007>

of morphogen interpretation itself from other self-organisational phenomena arising from cell-cell communication (i.e. spatially auto-correlated noise and cell-states); and therefore, to illustrate how the evolution of cell differentiation may benefit from learning-based principles to produce adaptive patterns under noisy conditions.

When individual cells experience explicit selection to perform such a function (complete information) during the training phase, they accurately recreate the complete segmented pattern in subsequent evolutionary stages provided they are exposed to all the required inputs. However, it can be the case that individual cells evolve in selective conditions that are missing information—i.e. where just a random subset of the multidimensional reaction norm is selected for and the remainder has not been subject to selection. In a developmental context that means that cells receive just a subset of the morphogen combinations they have evolved to recognise. Because of the intrinsic noisiness of many developmental systems, different cells in

each generation may receive different “bits” of incomplete positional information; the larger the amount of noise in the developmental system, the lesser the proportion of cells that will receive the whole set of correct signals during evolution [55,56]. From a theoretical perspective, the selective past of the cells’ GRNs can be viewed as a training phase, while the process of cell differentiation in subsequent generations can be viewed as the test phase. Normal development into an adult phenotype requires cells in the test phase to precisely interpret the morphogen inputs, even if the selective past has been noisy. Our experiment demonstrates that cells can overcome such noisy evolutionary past provided the developmental pattern (i.e. the adult phenotype) is built from a logically simple integration of the morphogens. That would imply that the logical structure of the multi-dimensional reaction norm that generates the segmented pattern is simple, so plastic cells can fit it from just a few points (that is, even if just a fraction of the cells have received the whole set of informative signals). Essentially, multi-dimensional reaction norms governed by simple logical functions belong to the category of problems that cells are well suited to solve because of their inductive bias towards simplicity (without this inductive bias, the response would be random and, in most cases, non-adaptive, S1 Fig). Thus, simple problems such as the one depicted in Fig 7A will be solved by the evolution of plastic cells better than at random. If, as here, the problem is to reproduce a specific developmental pattern, recognisable patterns will emerge even in very stochastic conditions, as Fig 7B shows.

In the same figure (lower row) we can also see how plastic cells fail to solve the same problem (and hence to rescue the corrupted pattern) if they do not have any inductive bias. In this case, cells are given *any* of the remaining logical functions compatible with the experienced ones, not just the simpler ones (which happen to be the adaptive ones in this example). For instance, consider a cell that has experienced two environments: In the first environment ($EF_1 = 1, EF_2 = 1$), the cell produced the phenotype $goi = -1$, whereas in the second environment ($EF_1 = 1, EF_2 = -1$), $goi = +1$. In the absence of any inductive bias, the cell will choose with equal probability any of the four logical functions compatible with these input/output relationships (XOR , $NAND$, $\neg EF_1$ and $EF_1 \leftrightarrow EF_2$). This randomness in the plastic response precludes cells from solving the plasticity problem presented (i.e. they are generally unable to repair corrupted phenotypes under noisy morphogen gradients, Fig 7B, S1 Fig).

Comparable noisy conditions are known to be commonplace in nature, and biological systems have developed a variety of mechanisms to cope with them [55, 56]. Our experiment suggests that under these non-ideal conditions, learning principles may explain natural selection’s ability to find successful buffering mechanisms. However, unlike other phenomena that increase developmental robustness by reducing the overall sensitivity of the system to the external conditions, such generalisation can increase developmental robustness while keeping the environmental sensitivity fully functional. This is possible because evolving GRNs represent simple *correlations* between morphogens rather than simply reducing the number or absolute concentrations of morphogens to which they are sensitive. In this way, an evolved GRN does not have to reduce the number of morphogens or the sensitivity of the system to these morphogens to be robust; instead its robustness is achieved by restricting the number of ways in which these morphogens can be combined to generate phenotypic outputs.

Discussion

In this work we have introduced a theoretical approach for the study of cell plasticity that exploits a formal mapping between cell plasticity and logical functions. Under this approach, the different multi-dimensional reaction norms exhibited by plastic cells are associated with specific logical (mathematical) functions derived from Boolean algebra. This idealisation enables us to measure the complexity of the different forms of cell plasticity.

By using this approach, we have shown that plastic cells are able to display complex forms of plasticity such as those associated by hard-to-compute, non-linearly decomposable logical functions ($\Omega = 1$). This class of logical functions is well known in learning theory because they cannot be learned by linear models (such as the linear perceptron [12]). To learn such functions requires the ability to represent and learn correlations between the inputs or to compare one input with another (e.g. inputs have different values, XOR), a feature which is mandatory for associative learning [12,14]. In computer science, the ability of artificial neural networks to represent and learn such functions is well understood, as it is in some biologically-inspired but network-free models of evolution [13]. In contrast, whether biologically realistic GRNs with an unstructured topology (e.g. an ad hoc topology without explicit hidden layers) and strictly non-negative variables ($[g_i] \geq 0$) can be evolved to represent such functions has not been previously shown. The implementation of learning theory approaches to study phenotypic evolution has been developed in prior work [14,15,26–29]. That work, however, is concerned with the evolution of evolvability (generating phenotypic variation given new genetic mutations, at least some of which may be adaptive) rather than adaptive plasticity (generating phenotypic variation given novel environmental cues, at least some of which may be adaptive). This paper thus focuses on how evolving systems exhibit learning principles when they classify a series of inputs (e.g. combinations of morphogens) into a set of discrete categories (e.g. cell states). In learning terms, the former corresponds to a generative model whereas the latter, addressed in this paper, corresponds to a discriminative model [12].

Experimental evidence from the field of synthetic biology shows that *in vivo* circuits (e.g. DNA-based logical gates or genetically engineered metabolic networks) are able to solve simple problems such as non-correlational AND-like functions ($\Omega \approx 0.6$). For more complex functions, the complexity of the biological circuits rapidly escalates with the problem complexity in a way that makes empirical approaches very challenging [57,58]. In those studies, the authors suggest that the lack of scalability is caused by the fact that real GRNs have to reuse the same genetic modules. In the present work, we have demonstrated that complex problems can be solved by biological systems if they are endowed with larger, denser and more strongly connected circuits (Fig 2). A network with signed states (e.g. [14]) can compute the same function using fewer nodes than a network with unsigned states (i.e., $[g_i] \geq 0$). This may explain why real GRNs, which can only have non-negative gene concentrations, must be large and densely connected in order to perform non-trivial computations [12].

Other computational models of plasticity avoid the complications created by non-negative states by following different strategies. Some introduce abstract state variables representing ‘gene expression potentials’, which are signed, rather than gene-expression levels, which are non-negative [14,15]. Others make use of an output layer to convert unsigned gene-expression states into signed phenotypic traits [38]. The most common approach is to directly encode the reaction norm parameters (e.g. slope) in genetically heritable variables, thus bypassing the generative dynamical processes responsible of the inductive biases [8]. Our results show that a conventional GRN model with unsigned states is sufficient to evolve such complex forms of adaptive plasticity.

Specifically, we demonstrated the potential of GRNs to generate complex multi-dimensional reaction norms under two biologically plausible scenarios: one in which the environmentally-sensitive factors are gene expression levels (classical GRNs) and one in which the environmentally-sensitive factors are the strengths of gene-gene interactions (second-order, tensor-based GRN). We show that the second-order implementation greatly improves the capacity of cells to evolve adaptive forms of plasticity, especially the more complex ones (i.e., those highly non-linear input-output maps). This observation suggests that more complex forms of cell plasticity can be expected in response to particular factors, such as temperature,

that influence expression dynamics through modulation of specific gene-gene interactions (not all interaction kinetics are equally sensitive to temperature changes) rather than individual gene-expression levels.

In addition, our results reveal that complex forms of plasticity are difficult to evolve because the hypervolumes of the GRN space containing them are very small and do not form a connected region (Fig 5B). The networks capable of complex plasticity are accordingly rare and require a dense topology of strong and heterogeneous connections. Likewise, the vast majority of new mutations in a given GRN drive the evolving cells into regions of equal or less complex plasticity. This applies to mutations that cause changes in GRN topology as well as mutations that cause changes in gene-gene interaction strengths [18]. By means of numerical simulations, we quantitatively predict the relative frequency of each type of plasticity from the probability of transitions between them (Fig 5B). This suggests that low evolvability of complex reaction norms is a consequence of the scarcity in the GRN space of the networks that generate complex reaction norms. A similar bias towards simple input-output maps is known for network-free evolutionary models [13] and has been recently proposed for artificial neural networks [23]. In addition, algorithmic theory suggests that it might be a general property of computable functions [44]. Our contribution is to demonstrate that this bias is intrinsic to cell plasticity, and that it emerges from reasonable model assumptions of gene-expression mechanisms. This may help to explain why many forms of plasticity tend to rely heavily on a low number of environmental cues even when many cues are potentially informative.

This bias towards simple forms of plasticity is consistent with studies on genotype-phenotype-maps (GPMs) which show that more complex phenotypes are more scattered in the parameter space and far less frequent than simple ones [17,37,59]. These qualitatively similar properties of GPMs and environment-phenotype-maps (EPMs) are likely due to the fact that they arise from the same underlying dynamical system: the GRN [37,43].

Although further work is required in order to systematically explore how complex GPMs and EPMs emerge spontaneously from GRN dynamics (and more complex multi-level developmental systems), our work provides a theoretical foundation for what to expect when plastic cells are exposed to novel multi-factorial environments. Knowing that cells preferentially acquire phenotypic states determined by the simplest combinations of the external factors may be useful for cell cultures exposed to new multi-nutrient growth media or for cancer research where tumor cells are treated with new multi-drug cocktails [20]. Although other models have addressed similar questions for specific cases (e.g. environmentally driven bacterial sporulation [16]), their approach excludes most GRN topologies, thus precluding the establishment of general principles.

In the context of cell plasticity, learning principles are relevant in three complementary ways. One is simply that by preferentially using simple logical functions, plastic (differentiating) cells do not equally consider all the available inputs (signalling molecules). Rather, cells ignore most inputs so that their phenotype depends on simple combinations of the remaining ones. This effectively decreases the size of the search space, making plastic cells evolve within a signalling environment that has an actual dimensionality much lower than the apparent one. In learning terminology, this is an inductive bias. Such biases favouring simple functions increase the chances of evolving an optimal combination of signals in a finite amount of time if the relationship between environmental cues and optimal phenotypes is, in fact, simple. For example, in the case of free-living single cells, if this bias is useful, it may enable cells to better track seasonal changes by using just a fraction of the information provided by the environment [31].

The second reason why learning principles are relevant is more subtle. Under the scenario we propose here, the differentiating cells preferentially establish simple logical relationships

between the environmental factors they have been exposed to (a form of inductive bias). This inductive bias towards simple functions means that plastic cells have an intrinsic, GRN-based mechanism for solving simple problems (although we have demonstrated that this bias is not limiting given the appropriate selective pressures). Thus, when the unforeseen challenge is a simple problem, the intrinsic bias towards simple plastic responses enables cells to show an improved performance compared to a random response (Fig 6A, Fig 7B, S1 Fig). This bias can be even more pronounced if the GRN connections are costly, a feature which is known to be common in real cells (Fig 6D). The difference to standard notions of Darwinian evolution is not trivial. In classical Darwinian evolution, plastic cells placed in a novel environment have no information to guess which will be the right (fitter) phenotype in that context [60]. In the absence of any inductive bias which causes cells to prefer one type of plasticity function over another, the cells will at best proceed by blind trial and error until the right cell state is eventually found, which may hinder or delay adaptation over evolutionary time [28].

The third reason why learning theory is relevant applies within a multicellular context: plastic cells that have evolved the cell differentiation function (which cell state corresponds to each combination of morphogens) can buffer the process of development against noisy or corrupted signalling pathways (Fig 7B). This is, in learning terms, an example of generalisation. Plastic cells achieve this by establishing only basic correlations between a few morphogens, which makes them consistently differentiate in the same phenotypic state (cell type) even when new signalling pathways arise by mutation. This inductive bias of the GRNs may increase phenotypic robustness against environmental or genetic perturbations, along with other well-known properties of biological GRNs like redundancy or modularity [7, 38]. Importantly, the learning-based robustness that we report here does not require a lack of sensitivity to the environmental inputs, but can arise from a logical simplification of the possible input-output relationships.

From a macroevolutionary perspective, it must be noted that the inductive bias towards simple forms of plasticity is present without being specifically selected for [61]. Rather, this bias towards simple plasticity is an inherent property of GRN dynamics which was already present in single-cell organisms well before the evolutionary transition towards multicellularity [5,6,61]. During this transition, which has happened at least 40 times [32,33], the functional integration of the new level of complexity crucially depends on the ability of the lower-level entities to correctly interpret and respond to a vastly rich set of inter-cellular signals [32]. Coming from free-living states, the multi-cellular state might have been very challenging [7]. However, as we have demonstrated here, the ability of cells to generalise from simple forms of plasticity may have facilitated the ability to adapt to more complex signalling environments. Testing how inherent inductive biases caused by cellular properties may have interacted with other factors in the self-organization of differentiated cell aggregates constitutes an exciting prospect for future research [62].

A possible criticism of our study is that plastic cells can only take advantage of adaptive generalisation if the plasticity function is simple (Fig 6A, S1 Fig). However, this does not undermine the generality of this phenomenon for two reasons. First, biologically, the signalling context inherent to multicellularity was not a pre-existing ecological niche to which cells adapted, but rather an environment that cells co-created themselves during the evolutionary transition [32,33]. From our findings, which show that single cells exhibit a bias favouring simple plasticity functions, it seems plausible that such simplicity pervaded the signalling context of the newly created multicellular environment. In other words, cells have evolved developmental patterning which requires only simple combinations of morphogens because this matches the simple plasticity they can exhibit. It is within this constructed space of limited dimensionality where plastic cells can take advantage of the learning principles; that is, cells

have created a problem that they were already able to solve. Second, logically, a bias that favours simple solutions fits well with generic properties of the natural world. This “universal logic”, familiar in probability theory, is often referred to as Occam’s razor, which states that all else being equal the simplest solution tends to be the correct one (the same reason explains why artificial NNs perform so well when confronted to complex problems [63]). The opposite bias (i.e., one favouring complex or less parsimonious solutions) would not produce effective learning in most cases. Notice that, since the bias towards simplicity is statistical (Fig 4A), it does not forbid the existence of some more complex than necessary solutions to evolutionary problems, which are well documented [64].

Another criticism may be that cells might have evolved mechanisms to overcome this bias towards simple forms of plasticity, for instance, by modifying the GRN connectivity. However, this would require an adaptive pressure towards denser GRNs, which *ceteris paribus* seems unlikely given metabolic costs of regulation [15,65]. A more plausible scenario, that GRN connections are costly, results in an even more pronounced bias towards simple forms of plasticity (Fig 6D), so this bias seems to be natural.

Intuitively, the ability of plastic cells to mimic logical functions and to perform learning with generalisation is to be expected given the structural, dynamical, and functional isomorphisms between evolving GRNs and learning NNs: they both partition the space of possible inputs into a series of states or attractors [17,41,59], both are adaptable via incremental improvement principles [19,37,66] and both exhibit similar inductive biases [18, 44]. In addition, both exhibit increased performance under similar adaptive pressures [15]. For example, we have shown that some procedures commonly used to improve the performance of artificial NNs (such as the L_1 and L_2 regularisation) can also enhance the ability of plastic cells to acquire specific reaction norms from incomplete information (Fig 6D).

These structural and functional commonalities between GRNs and NNs enable the transfer of knowledge derived from learning theory to illuminate the domain of cell plasticity. This approach is possible because learning principles and logical rules are not substrate specific: they can be exhibited by interacting molecules, genes, cells, neurons and transistors alike as long as these elements have the right type of interactions. This substrate independence of learning principles helps interpret the main finding of this work, namely, that evolving gene networks can exhibit similar adaptive principles to those already familiar in cognitive systems.

Methods

Evolutionary protocol

Many of our experiments involve the adaptive evolution of the GRNs towards different multi-dimensional reaction norms (that is, to match specific input-output maps or truth tables), including those described by complex logical functions (Fig 1) [39]. Whilst some authors have implemented this as a single population which faces sequentially all possible environments over a number of generations (each environment being a specific input-output row of the whole truth table [15,19,30]), we split in each generation the entire population of plastic cells into a set of minimal (single-cell) isogenic sub-populations, with each sub-population being exposed to a specific environment (i.e. a specific combination of the environmental factors). This procedure is equivalent to the first one but allow us to save computational resources [12] and to ensure that all environments equally contribute to the acquisitions of specific forms of cell plasticity.

In each generation, the target phenotype P^* (an arbitrary function of known complexity) is compared to the resulting phenotype P (each element of P is the binary expression of *goi* in each of the possible environments). The sum of all the matches between the phenotype P and

the target P^* , scaled by the size of phenotypic vector, determines the fitness W of a given GRN. The more similar the phenotype to the optimal (target) one, the higher the fitness W . The fitness W determines the likelihood of a GRN to contribute to the next generation:

$$W = \frac{\sum_{i=1}^{2^{Efs}} |P_i - P_{*i}|}{2^{Efs}} \quad (11)$$

This is a fitness measure based on the distance between P and P^* , which is widely used as a proxy of biological aptitude [36,37,59]. We use Hamming distances instead of Euclidean distances because our phenotypes have binary expression. In each generation, the alleles encoding for the GRN undergo point mutations: in a given generation, each element B_{ij} changes its value to $B_{ij} + v$ with a probability $p = 1/Ng$, being $v \sim U(-0.1, 0.1)$, which ensures that mutational effects are independent of the GRN size. We do not mutate the G vector, which is initialized to the same initial value as in $t = 0$ in each generation (we keep it constant because we are interested in the environmental-phenotype properties of the system, rather than in those properties derived from the more widely studied genotype-phenotype-maps *sensu* [17]). Each mutational event uniquely determines a mutant phenotype, whose probability of being fixed is proportional to the selective advantage of that mutation. Thus, we confront again all the clone subpopulations bearing the mutant allele with all possible environments, recording the mutant's fitness Wm . If $Wm \geq W$, then the new mutation is fixed.

We use this hill-climbing evolutionary algorithm because our aim is to provide a phenomenological model of the evolution of structured phenotypic plasticity, rather than quantifying realistic evolutionary rates or other important population or genetic parameters (e.g. magnitude of selection coefficients, mutation probabilities, population structure, epistatic effects or the evolution of polymorphism) [25]. In doing so, we assume that the selection coefficient is sufficiently large, and that mutation rate is sufficiently low, that beneficial mutations get deterministically fixed and deleterious mutations do not [34]. This represents the simplest algorithm for adaptive evolution [67], and allows us to demonstrate that more stochastic evolutionary scenarios with different dynamic properties, including those allowing for the coexistence of GRN polymorphisms, are not required to exhibit the learning capabilities we study. The implications of relaxing this limitation remains for further research.

Generalisation experiments

In experiment 1, natural selection can find arbitrary forms of cell plasticity, but this required cells to be exposed to all possible combinations of environmental factors (complete information). In the experiments 4 and 4b, plastic cells are exposed to a subset of all possible environments (referred as the training set Ts ($Ts < Ne$); incomplete information). We evolve the system in a training phase until the maximum possible fitness under that specific Ts size (so that cells produce the required phenotypic states for each of the Ts environments). For each replicate, the composition of the training set is randomly chosen from the set of 2^{Ne} possible environments.

In the second (test) phase of the experiment, we expose the trained GRNs to the remaining $2^{Ne} - Ts$ environments, which are utterly novel to the cells. The fitness W is calculated now over all the 2^{Ne} environments. If the cells are able to generalize from their past evolutionary experience, the performance of the plastic cells in these new environments will be better than at random: $W \geq (2^{Ne} - Ts)/2$ (Fig 6A and 6B, and S1 Fig, green lines).

In the experiment 4b, where generalisation occurs in an explicit spatial 2D field of 1250 cells (25x50 sub-hexagonal grid) instead of a space-less implementation (Fig 7), the target developmental pattern was chosen amongst a large set of random patterns resulting from the

integration of three morphogens (each one can be viewed as a spatially distributed environmental factor) according to random thresholds. The pattern, resembling an arthropodian early embryo, was specifically chosen to be generated from simple logical functions and visually recognizable as a biological structure. The spatial distributions (xy coordinates) of the morphogens were drawn from simple and biologically realistic gradients: a Turing-like striped pattern ($EF_1(x,y) = \sin(x \cdot k)$), an isotropic radial gradient from a punctual signalling centre ($EF_2(x,y) = (x^2 + y^2)^{1/2}$) and a lateral diffusion from the antero-posterior midline ($EF_3(x,y) = y^{1/2}$)

The combination of these morphogens generates the target developmental pattern according to the composite logical function $AND(OR(AND(EF_1,EF_2),AND(EF_2,EF_3)),EF_1)$ which has a relatively low complexity $\Omega \approx 0.5$. This function is associated with a truth table of length $2^{Ne} = 32$ (an arbitrary threshold was applied to the continuous morphogen gradients in order to reduce the length of the truth table). For the training phase, each individual cell was exposed to a random subset of this table ($TS < 32$). For the test phase, every cell is exposed to the whole information ($TS = 32$). We contrast these results with another scenario in which the missing information in the test phase is generated randomly (Fig 7B, S1 Fig).

Supporting information

S1 Fig. Developmental implications of learning-based cell plasticity. Associated learning-TS curves for experiment 4B (Fig 7). In this experiment, cells are required to integrate several morphogen gradients in a phenotypic response arising from the combination of simple logical functions (see Methods). Since the structure of the problem coincides with the inductive bias plastic cells have (a preference for simple forms of cell plasticity), they perform better than at random when some information is missing. To simulate this scenario of missing information, we make randomly chosen cells to experience only a random subset of the truth table (the whole table contains the whole sequence of logical steps shown in Fig 7A, which are necessary to recreate the segmented developmental pattern, see Methods). Red line represents the information provided ($x = y$), green line represents the expected performance at random $(2^{Ne} - TS) / 2$, lower row of images in Fig 7B; and blue line the degree of matching between the resulting phenotypes and the expected ones (the ones which generate the striped developmental pattern; upper row in Fig 7B). TS is the size of the training set. $Ne = 5$; $n = 30$ replicates, standard deviation as boxes, extreme values as error bars. The conclusion drawn from the figure is that thanks to learning principles plastic cells are able to restore corrupted phenotypes arising from simple combinations of exogenous signals.

(EPS)

S2 Fig. Effects of including costly GRN connections in the evolutionary algorithm. Points show the averages over $n = 30$ replicates for an environmental dimensionality of $Ne = 3$. Target plastic functions have either high logical complexity ($\Omega \approx 1.0$, red lines) or low complexity ($\Omega < 0.5$, blue lines). The metabolic cost of GRN connections is implemented by two different regularization procedures taken from computer sciences (See main text for a biological and methodological rationale of these procedures): L_1 and L_2 regularization, depicted by solid and dashed lines respectively. The X axis is the relative weight of the costly connections in determining the individual fitness (parameter in the model, see Experiment 4a; Fig 6D).

(EPS)

Acknowledgments

The authors thank T. Uller, A. Rago, D. Prosser, J. Caldwell, J.F. Nash, L.E. Mears and I. Herando-Herraez for helpful discussion and comments.

Author Contributions

Conceptualization: Miguel Brun-Usan, Richard A. Watson.

Formal analysis: Miguel Brun-Usan, Christoph Thies, Richard A. Watson.

Funding acquisition: Richard A. Watson.

Investigation: Miguel Brun-Usan, Christoph Thies, Richard A. Watson.

Methodology: Christoph Thies, Richard A. Watson.

Project administration: Richard A. Watson.

Software: Miguel Brun-Usan.

Supervision: Richard A. Watson.

Visualization: Miguel Brun-Usan.

Writing – original draft: Miguel Brun-Usan, Richard A. Watson.

Writing – review & editing: Christoph Thies, Richard A. Watson.

References

1. Kropotkin P. Mutual aid: a factor in evolution. London, UK. Heinemann. 1902.
2. West-Eberhard MJ. Developmental plasticity and evolution. Oxford University Press. 2003.
3. Laland KN, Uller T, Feldman MW, Sterelny K, Müller GB, Moczek A, et al. The extended evolutionary synthesis: its structure, assumptions and predictions. *Proc R Soc B*. 2015; 282(1813): 20151019. <https://doi.org/10.1098/rspb.2015.1019> PMID: 26246559
4. Gilbert SF. The morphogenesis of evolutionary developmental biology. *Int J Dev Biol*. 2003; 47:467–77. PMID: 14756322
5. Tudge SJ, Watson RA, Brede M. Game theoretic treatments for the differentiation of functional roles in the transition to multicellularity. *J Theor Biol*. 2016; 395:161–73. <https://doi.org/10.1016/j.jtbi.2016.01.041> PMID: 26869214
6. Sebé-Pedrós A, Degnan BM, Ruiz-Trillo I. The origin of Metazoa: a unicellular perspective. *Nat Rev Genet*. 2017; 18:498–512. <https://doi.org/10.1038/nrg.2017.21> PMID: 28479598
7. Gilbert SF, Raunio AM. Embryology: constructing the organism. Sinauer. Sunderland, MA. 1997.
8. De Jong G. Evolution of phenotypic plasticity: patterns of plasticity and the emergence of ecotypes. *New Phytol*. 2005; 166(1):101–18. <https://doi.org/10.1111/j.1469-8137.2005.01322.x> PMID: 15760355
9. Griffiths AJ, Wessler SR, Lewontin RC, Gelbart WM, Suzuki DT, Miller JH. An introduction to genetic analysis. Macmillan. 2005.
10. Braun E. The unforeseen challenge: from genotype-to-phenotype in cell populations. *Rep Prog Phys*. 2015; 78:036602. <https://doi.org/10.1088/0034-4885/78/3/036602> PMID: 25719211
11. Munro E, Robin F, Lemaire P. Cellular morphogenesis in ascidians: how to shape a simple tadpole. *Curr Op Gen Dev*. 2006; 16:399–405.
12. Sipser M. Introduction to the theory of computation. Vol. 2. Boston: Thomson Course Technology. 2006.
13. Lenski RE, Ofria C, Pennock RT, Adami C. The evolutionary origin of complex features. *Nature*. 2003; 423(6936):139. <https://doi.org/10.1038/nature01568> PMID: 12736677
14. Watson RA, Wagner GP, Pavlicev M, Weinreich DM, Mills R. The evolution of phenotypic correlations and “developmental memory”. *Evolution*. 2014; 68(4):1124–38. <https://doi.org/10.1111/evo.12337> PMID: 24351058
15. 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 Comp Biol*. 2017; 13(4):e1005358.
16. van Gestel J, Weissing FJ. Regulatory mechanisms link phenotypic plasticity to evolvability. *Sci Rep*. 2016; 6:24524. <https://doi.org/10.1038/srep24524> PMID: 27087393
17. Alberch P. From genes to phenotype: dynamical systems and evolvability. *Genetica*. 1994; 84(1):5–11.

18. Salazar-Ciudad I. Developmental constraints vs variational properties: how pattern formation can help to understand evolution and development. *J Exp Zool B Mol Dev Evol*. 2006; 306(2):107–25. <https://doi.org/10.1002/jez.b.21078> PMID: 16254986
19. Crombach A, Hogeweg P. Evolution of evolvability in gene regulatory networks. *PLoS Comp Biol*. 2008; 4(7): e1000112.
20. Huang S, Eichler G, Bar-Yam Y, Ingber DE. Cell fates as high-dimensional attractor states of a complex gene regulatory network. *Phys Rev Lett*. 2005; 94:128701. <https://doi.org/10.1103/PhysRevLett.94.128701> PMID: 15903968
21. Wang J, Zhang K, Xu L, Wang E. Quantifying the Waddington landscape and biological paths for development and differentiation. *Proc Natl Acad Sci USA*. 2011; 108:8257–62. <https://doi.org/10.1073/pnas.1017017108> PMID: 21536909
22. Font-Clos F, Zapperi S, La Porta CA. Topography of epithelial-mesenchymal plasticity. *Proc Natl Acad Sci USA*. 2018; 115:5902–7. <https://doi.org/10.1073/pnas.1722609115> PMID: 29784817
23. Pérez GV, Louis AA, Camargo CQ. Deep learning generalizes because the parameter-function map is biased towards simple functions; 2018. Preprint. Available from: arXiv:1805.08522. Cited 16 October 2018.
24. Abraham A. Artificial neural networks. Sydenham P Thorn R, eds. In: *Handbook of Measuring System Design*. London: John Wiley & Sons Inc, 2005.
25. Kashtan N, Noor E, Alon U. Varying environments can speed up evolution. *Proc Natl Acad Sci USA*. 2007; 104(34):13711–6. <https://doi.org/10.1073/pnas.0611630104> PMID: 17698964
26. Watson RA, Szathmáry E. How can evolution learn?. *Trends Ecol Evol*. 2016; 31(2):147–57. <https://doi.org/10.1016/j.tree.2015.11.009> PMID: 26705684
27. Chastain E, Livnat A, Papadimitriou C, Vazirani U. Algorithms, games, and evolution. *Proc Natl Acad Sci USA*. 2014; 111(29):10620–3. <https://doi.org/10.1073/pnas.1406556111> PMID: 24979793
28. Mayley G. Landscapes, learning costs, and genetic assimilation. *Evol Comput*. 1996; 4(3):213–34.
29. Power DA, Watson RA, Szathmáry E, Mills R, Powers ST, Doncaster CP, et al. What can ecosystems learn? Expanding evolutionary ecology with learning theory. *Biol Direct*. 2015; 10(1):69.
30. Xue BK, Leibler S. Evolutionary learning of adaptation to varying environments through a transgenerational feedback. *Proc Natl Acad Sci USA*. 2016; 113(40):11266–71. <https://doi.org/10.1073/pnas.1608756113> PMID: 27647895
31. Benincà E, Huisman J, Heerkloss R, Johnk KD, Branco P, Van Nes EH, et al. Chaos in a long-term experiment with a plankton community. *Nature*. 2008; 451(7180):822–5. <https://doi.org/10.1038/nature06512> PMID: 18273017
32. Maynard-Smith J, Szathmáry E. *The major transitions in evolution*. Freeman. NY. 1995.
33. Grosberg RK, Strathman RR. The evolution of multicellularity: a minor major transition? *Ann Rev Ecol Evol Syst*. 2007; 38(1):621–54.
34. Pavlicev M, Cheverud JM, Wagner GP. Evolution of adaptive phenotypic variation patterns by direct selection for evolvability. *Proc Roy Soc B Biol Sci*. 2011; 278:1903–12.
35. Cornish-Bowden A. *Fundamentals of enzyme kinetics*. London. Butterworths. 1979.
36. Salazar-Ciudad I, Garcia-Fernandez J, Solé RV. Gene networks capable of pattern formation: from induction to reaction–diffusion. *J Theor Biol*. 2000; 205(4):587–603. <https://doi.org/10.1006/jtbi.2000.2092> PMID: 10931754
37. Salazar-Ciudad I, Newman SA, Solé RV. Phenotypic and dynamical transitions in model genetic networks I. Emergence of patterns and genotype-phenotype relationships." *Evol Dev*. 2001; 3(2):84–94. <https://doi.org/10.1046/j.1525-142x.2001.003002084.x> PMID: 11341677
38. Draghi JA, Whitlock MC. Phenotypic plasticity facilitates mutational variance, genetic variance, and evolvability along the major axis of environmental variation. *Evolution*. 2012; 66(9):2891–902. <https://doi.org/10.1111/j.1558-5646.2012.01649.x> PMID: 22946810
39. Wagner A. Does evolutionary plasticity evolve? *Evolution*. 1996; 50(3):1008–23. <https://doi.org/10.1111/j.1558-5646.1996.tb02342.x> PMID: 28565284
40. Pavlicev M, Hansen TF. Genotype-phenotype maps maximizing evolvability: modularity revisited. *Evol Biol*. 2011; 38(4):371–89.
41. Kauffman SA. Metabolic stability and epigenesis in randomly constructed genetic nets. *J Theor Biol*. 1969; 22:437–67. [https://doi.org/10.1016/0022-5193\(69\)90015-0](https://doi.org/10.1016/0022-5193(69)90015-0) PMID: 5803332
42. Mengistu H, Huizinga J, Mouret JB, Clune J. The evolutionary origins of hierarchy. *PLoS Comp Biol*. 2016; 12(6):e1004829.

43. Furusawa C, Kaneko K. Formation of dominant mode of evolution in biological systems. *Phys Rev E*. 2018; 97(4):042410.
44. Dingle K, Camargo CQ, Louis AA. Input-output maps are strongly biased towards simple outputs. *Nat Comm*. 2018; 9(1):761.
45. Wibral M, Priesemann V, Kay JW, Lizier JT, Phillips WA. Partial information decomposition as a unified approach to the specification of neural goal functions. *Brain Cogn*. 2017; 112:25–38. <https://doi.org/10.1016/j.bandc.2015.09.004> PMID: 26475739
46. Ebeling W, Jiménez-Montaño MA. On grammars, complexity, and information measures of biological macromolecules. *Math Biosc*. 1980; 52:53–71.
47. Bennet CH. Logical depth and physical complexity. In: Herken R. (Ed.) *The Universal Turing Machine: A half Century of Survey*. Vienna: Springer-Verlag. 1995
48. MacKay DJC. *Information theory, inference and learning algorithms*. Cambridge university press. 2003.
49. Dunker AK, Brown CJ, Lawson JD, Iakoucheva LM, Obradovic Z. Intrinsic disorder and protein function. *Biochemistry*. 2002; 41(21):6573–82. <https://doi.org/10.1021/bi012159+> PMID: 12022860
50. Wright PE, Dyson HJ. Intrinsically disordered proteins in cellular signalling and regulation. *Nature Rev Mol Cell Biol*. 2015; 16:18.
51. Niklas KJ, Bondos SE, Dunker AK, Newman SA. Rethinking gene regulatory networks in light of alternative splicing, intrinsically disordered protein domains, and post-translational modifications. *Front Cell Dev Biol*. 2015; 3(8):1–13.
52. Kaneko K, Furusawa C, Yomo T. Universal relationships in gene-expression changes for cells in steady-growth state. *Phys rev X*. 2015; 5(1):011014
53. Gregor T, McGregor AP, Wieschaus EF. Shape and function of the Bicoid morphogen gradient in dipteran species with different sized embryos. *Dev Biol*. 2008; 316(2):350–8. <https://doi.org/10.1016/j.ydbio.2008.01.039> PMID: 18328473
54. Jaeger J. The gap gene network. *Cell Mol Life Sci*. 2011; 68(2):243–74. <https://doi.org/10.1007/s00018-010-0536-y> PMID: 20927566
55. Kaneko K. Evolution of robustness to noise and mutation in gene expression dynamics. *PloS One* 2007; 2(5):e434. <https://doi.org/10.1371/journal.pone.0000434> PMID: 17502916
56. Eldar A, Elowitz MB. Functional roles for noise in genetic circuits. *Nature* 2010; 467(7312):167–73. <https://doi.org/10.1038/nature09326> PMID: 20829787
57. Adleman LM. Molecular computation of solutions to combinatorial problems. *Science*. 1994; 206:1021–4.
58. Macia J, Posas F, Solé RV. Distributed computation: the new wave of synthetic biology devices. *Trends Biotechnol*. 2012; 30(6):342–9. <https://doi.org/10.1016/j.tibtech.2012.03.006> PMID: 22516742
59. Salazar-Ciudad I, Marin-Riera M. Adaptive dynamics under development-based genotype-phenotype maps. *Nature*. 2013; 497:361–4. <https://doi.org/10.1038/nature12142> PMID: 23636325
60. Ghalambor CK, McKay JK, Carroll SP, Reznick DN. Adaptive versus non-adaptive phenotypic plasticity and the potential for contemporary adaptation in novel environments. *Funct Ecol*. 2007; 21:394–407.
61. Newman SA, Forgacs G, Muller GB. Before programs: the physical origination of multicellular forms. *Int J Dev Biol*. 2003; 50(2–3):289–99.
62. Solé RV, Valverde S. Before the endless forms: embodied model of transition from single cells to aggregates to ecosystem engineering. *PloS one*. 2013; 8(4):e59664. <https://doi.org/10.1371/journal.pone.0059664> PMID: 23596506
63. Lin HW, Tegmark M, Rolnick D. Why does deep and cheap learning work so well?. *J Stat Phys*. 2017; 168(6):1223–47.
64. Westerhoff HV, Winder C, Messiha H, Simeonidis E, Adamczyk M, Verma M, Bruggeman FJ, Dunn W. Systems biology: the elements and principles of life. *FEBS letters*. 2009; 583(24):3882–90. <https://doi.org/10.1016/j.febslet.2009.11.018> PMID: 19913018
65. Clune J, Mouret JB, Lipson H. The evolutionary origins of modularity. *Proc Roy Soc B*. 2013; 280(1755):20122863.
66. Gerhart J, Kirschner M. *Cells, embryos and evolution: toward a cellular and developmental understanding of phenotypic variation and evolutionary adaptability*. Blackwell Science. Malden, MA. 1997.
67. Nowak MA. *Evolutionary dynamics*. Harvard university press. 2006.