

Research



Cite this article: Catalán P, Manrubia S, Cuesta JA. 2020 Populations of genetic circuits are unable to find the fittest solution in a multilevel genotype–phenotype map. *J. R. Soc. Interface* **17**: 20190843. <http://dx.doi.org/10.1098/rsif.2019.0843>

Received: 11 December 2019

Accepted: 12 May 2020

Subject Category:

Life Sciences–Mathematics interface

Subject Areas:

computational biology, evolution, systems biology

Keywords:

genotype–phenotype map, toyLIFE, gene regulatory networks, phenotypic bias, entropy, genetic circuits

Author for correspondence:

Pablo Catalán

e-mail: pablocatalanfdz@gmail.com

Electronic supplementary material is available online at <https://doi.org/10.6084/m9.figshare.c.4987157>.

Populations of genetic circuits are unable to find the fittest solution in a multilevel genotype–phenotype map

Pablo Catalán^{1,2}, Susanna Manrubia^{1,3} and José A. Cuesta^{1,2,4,5}

¹Grupo Interdisciplinar de Sistemas Complejos (GISC), Madrid, Spain

²Departamento de Matemáticas, Universidad Carlos III de Madrid, Leganés, Madrid, Spain

³Departamento de Biología de Sistemas, Centro Nacional de Biotecnología (CSIC), Madrid, Spain

⁴Instituto de Biocomputación y Física de Sistemas Complejos (BIFI), Universidad de Zaragoza, Zaragoza, Spain

⁵UC3M-Santander Big Data Institute (IBiDat), Universidad Carlos III de Madrid, Getafe, Madrid, Spain

PC, 0000-0003-2826-4684; SM, 0000-0003-0134-2785; JAC, 0000-0001-9890-9367

The evolution of gene regulatory networks (GRNs) is of great relevance for both evolutionary and synthetic biology. Understanding the relationship between GRN structure and its function can allow us to understand the selective pressures that have shaped a given circuit. This is especially relevant when considering spatio-temporal expression patterns, where GRN models have been shown to be extremely robust and evolvable. However, previous models that studied GRN evolution did not include the evolution of protein and genetic elements that underlie GRN architecture. Here we use *toyLIFE*, a multilevel genotype–phenotype map, to show that not all GRNs are equally likely in genotype space and that evolution is biased to find the most common GRNs. *toyLIFE* rules create Boolean GRNs that, embedded in a one-dimensional tissue, develop a variety of spatio-temporal gene expression patterns. Populations of *toyLIFE* organisms choose the most common GRN out of a set of equally fit alternatives and, most importantly, fail to find a target pattern when it is very rare in genotype space. Indeed, we show that the probability of finding the fittest phenotype increases dramatically with its abundance in genotype space. This phenotypic bias represents a mechanism that can prevent the fixation in the population of the fittest phenotype, one that is inherent to the structure of genotype space and the genotype–phenotype map.

1. Introduction

The evolution of gene regulatory networks (GRNs) is a topic of great relevance [1,2]. Organisms show a plethora of complex regulatory architectures in order to carry out several developmental programmes [3] and to integrate signals from the environment [4]. As a result, much work has been devoted to understanding how these architectures have evolved, and to disentangling the relationship between the structure of a GRN and its function [5,6]. The underlying motivation is to understand which regulatory motifs appear as a result of selection for a given function or, conversely, what kind of functionality is attained when the structure of the GRNs is determined by other factors. GRNs are also the object of intense research from the standpoint of synthetic biology, which tries to design circuits to perform pre-defined functions [7].

One of regulation's most interesting outcomes is the generation of spatio-temporal patterns of gene expression that multicellular organisms use in their development [8]. Recent work has been devoted to the study of the architecture of GRNs that give rise to different patterns, exploring their robustness and evolvability [9–12]. These studies have found that GRNs can easily evolve to generate new patterns, facilitating the emergence of new developmental programmes. The same pattern can be achieved by means of very different mechanisms [11], which in turn determine the levels of robustness and

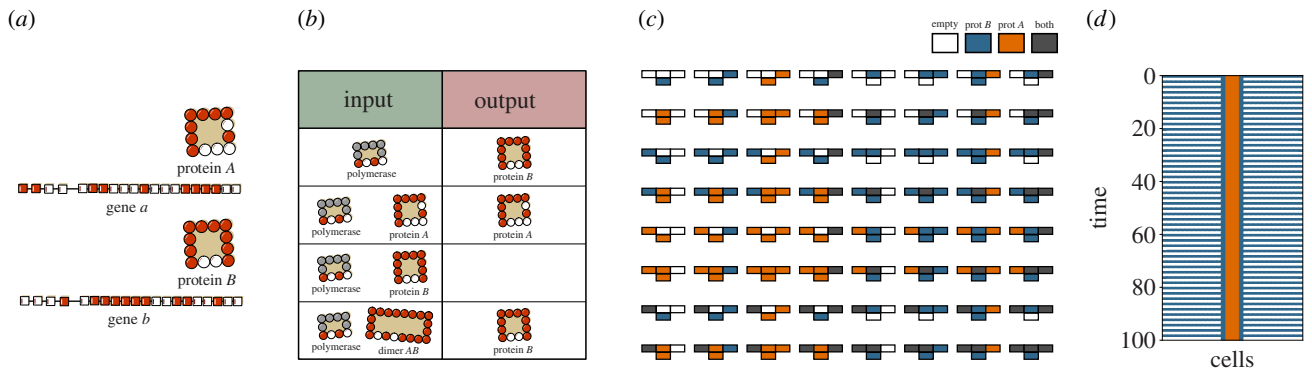


Figure 1. toyLIFE is a multilevel genotype–phenotype map. (a) toyLIFE genotypes are binary strings of length $20n$, where n is the number of genes in the genome. The first four letters of each gene represent its promoter region, while the remaining 16 are the coding region. The coding region, when expressed, turns into a protein that folds into a 4×4 lattice (see Supplementary text in the electronic supplementary material). (b) Following toyLIFE 's interaction rules, we obtain the corresponding gene regulatory network (GRN), represented here by its truth table. (c) Each GRN determines, under some propagation rules, a unique cellular automaton. Given the state of a cell and its neighbours at time t , toyLIFE 's rules determine the state of the cell at time $t + 1$, where cells can be empty (white), expressing protein A (orange), expressing protein B (blue) and expressing both proteins (grey). (d) Under certain initial conditions (in this case, the expression of protein A in the middle cell of the tissue), the cellular automata give rise to spatio-temporal patterns of gene expression. In this case, the cellular automaton in c leads to an alternating pattern in which the tissue expresses protein B and then does not express anything, while in the centre of the tissue three cells express protein A continuously.

evolvability of the pattern. However, GRNs are the result of interactions between proteins and genetic elements, and the evolution of GRNs is a direct result of changes in protein folding, binding affinities and promoter or enhancer regions. Owing to its enormous complexity, models of GRN evolution rarely incorporate these underlying dynamics, although there are some exceptions [13,14].

Here we use a multilevel computational model of gene regulation to show that some GRN architectures are easier to build from interacting proteins and genes than others. As a result, there is a phenotypic bias [15–17] that turns some GRNs into attractors of evolutionary dynamics, even in the absence of fitness differences.

We focus on Boolean GRNs, in which genes can either be ON or OFF [18,19]. Although different from other models of gene expression, where the concentration of proteins can vary continuously [20], Boolean networks have been repeatedly used to model GRN evolution [21,22], and some regulatory functions have been explained best by using Boolean functions [23]. Our Boolean GRNs are also modelled in discrete time, so that the expression of one cell in time $t + 1$ is determined by its expression and that of its neighbouring cells in time t . This formalism transforms GRNs into cellular automata [24]. Connecting several cells in a one-dimensional tissue, and allowing for propagation of gene products between neighbouring cells, we obtain spatio-temporal patterns that are similar to those found in real organisms.

These Boolean GRNs are built on top of a simple model of cellular biology, toyLIFE [25,26]. toyLIFE organisms contain genes, which are translated into proteins that interact with each other to form dimers. Both dimers and proteins alter the expression of genes, thus creating Boolean GRNs such as those described above. As a consequence, toyLIFE is a multilevel map from binary genomes (genotypes) to Boolean GRNs (first phenotype level) to cellular automata (second phenotype level) to spatio-temporal patterns (third phenotype level) (figure 1), thus allowing us to study the effects of molecular evolution at different phenotypic levels.

We show that toyLIFE genomes with two genes are able to generate a wide variety of GRNs and spatio-temporal patterns. Moreover, not all of these are equally abundant in

Table 1. Truth table for a two-gene GRN.

$A(t)$	$B(t)$	$A(t + 1)$	$B(t + 1)$
0	0	1	0
0	1	0	1
1	0	1	0
1	1	0	0

genotype space: some GRNs are mapped by many genotypes, while others are comparatively rare. We find that this phenotypic bias is enough to steer evolving populations towards more abundant GRNs, thus introducing an additional element when trying to explain GRN evolution: one that is not related to function or structure. Furthermore, we also show that this phenomenon can result in the inability of the evolutionary search to find some regulatory patterns, even when they are fitter than every other.

2. Results

2.1. Boolean networks and spatio-temporal patterns

In a Boolean GRN, a gene can either be ON (expressed) or OFF (not expressed). The expression state of each gene at time $t + 1$ is a function of the expression states of all the other genes in the network at time t , so that each state of the network maps into another state. In a GRN with two genes a and b and corresponding proteins A and B , this mapping can be represented as a truth table that connects every input state to an output state. Table 1 shows an example of this: the state (0, 0) is mapped to (1, 0), which means that gene a is expressed constitutively. The next two rows indicate that both a and b activate their own expression, while the last row shows that both genes repress each other. The truth table determines the temporal expression patterns of a Boolean GRN, thus giving us all the information we need to study this system.

tissue: one-dimensional row formed by 31 cells

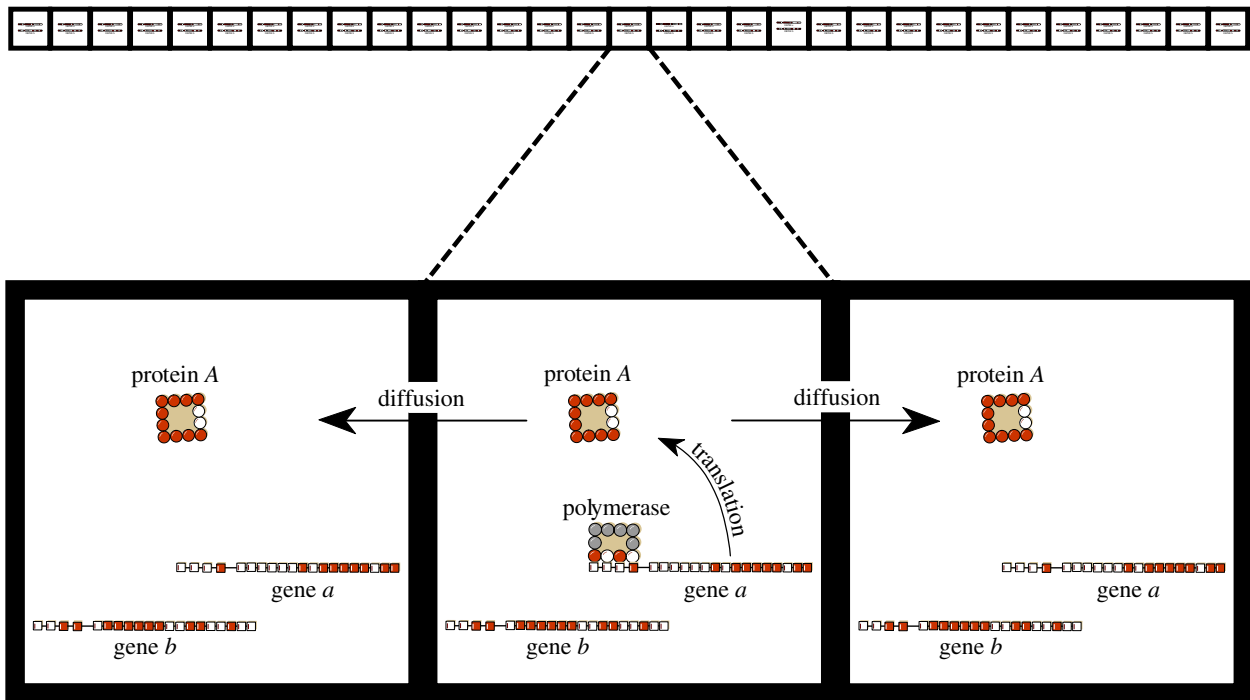


Figure 2. Pattern-formation phenotype in toyLIFE . We consider a one-dimensional row formed by 31 cells. The figure illustrates one example of how this multicellular phenotype works using toyLIFE for illustration purposes. When protein A is expressed, it can propagate to neighbouring cells and influence gene expression there. This way, the spatio-temporal state of the tissue becomes a cellular automaton.

We want to study the spatio-temporal patterns of two-gene GRNs embedded in a one-dimensional tissue. First, we define the number of cells in the tissue, which we will consider to be constant. For our purposes, we choose tissues with 31 cells in a row. The number of cells is arbitrary and it does not affect our results: the same patterns are generated by the same truth tables under similar regulatory inputs (electronic supplementary material, figure S8), so no phenomenology is lost from restraining our study to this tissue size.

Now we define the connections between different cells in the tissue. We will assume that only protein A can propagate to the adjoining cells (figure 2). As a result, the input state of cell c_i in time $t + 1$ will be affected by the output states of cells c_{i-1} and c_{i+1} in time t —as well as its own. We will further assume that there is enough protein A to stay inside the cell and propagate to the adjoining ones. For the cells at the beginning and end of the tissue, we impose the following boundary condition: cell c_1 will be affected by itself and cell c_2 , and cell c_L (where L is the length of the tissue) will be affected by itself and c_{L-1} —remember that $L = 31$ throughout.

With these rules, each GRN (defined by a truth table) gives rise to a cellular automaton [24] (figure 1c) with four states: (0) no protein is expressed (white), (1) protein B is expressed (blue), (2) protein A is expressed (orange), and (3) both proteins are expressed (grey) (see electronic supplementary material, figure S9 to see how the cellular automaton update rule is obtained from the truth table). Cellular automata are compactly described by the output they produce given an input. Because the input of a cell is formed by itself and its adjoining cells, and because each of them can be in four states, the number of input states is $4^3 = 64$. The number of possible cellular automata is, therefore, $4^{64} \approx 3.4 \times 10^{38}$. We will see below that the number of two-gene toyLIFE genotypes, which we use to generate our GRNs, is around 10^{12} , which is not enough to explore this vast number.

Each cellular automaton, in turn, gives rise to a spatio-temporal pattern that will depend on the initial conditions of the tissue at time $t = 0$. It is soon evident that the space of these scenarios is hyper-astronomical in size [17], and so we choose to start our dynamics with both genes in every cell in the OFF state, except the cell in the middle of the tissue (c_{15}), where we will express protein A, modelling a signal received from the exterior of the tissue. We then explore the expression dynamics of the whole tissue for 100 time steps, enough to resolve all patterns.

We now explore four relevant GRNs and their resulting patterns (figure 3). They are (a) a double-negative feedback loop with self-activation, (b) the same as before but with gene a having a constitutive expression, (c) a double-positive feedback loop with self-activation, and (d) a double-positive feedback loop without self-activation. Figure 3 shows the truth tables associated with these GRNs and the patterns they generate under the conditions mentioned above. The first two patterns result in protein A being expressed in a stable manner in the whole tissue. The difference between them is that in figure 3b protein A is expressed constitutively in every cell, while in figure 3a that signal must propagate through the tissue. The pattern in figure 3c is similar to the one in figure 3a, but both proteins end up expressed in the tissue, as a result of the positive feedback loop. Finally, in figure 3d the tissue expresses protein A and B in an alternating way.

Let us focus on the pattern generated by the network in figure 3b. There are 16 GRNs that generate the same pattern under the conditions defined above (electronic supplementary material, figure S10 shows the truth tables for all of these). If there were selection pressures to create that particular pattern, we could expect evolutionary dynamics to choose among these 16 GRNs with equal probability, everything else being equal. This is certainly what almost every mathematical model of phenotypic

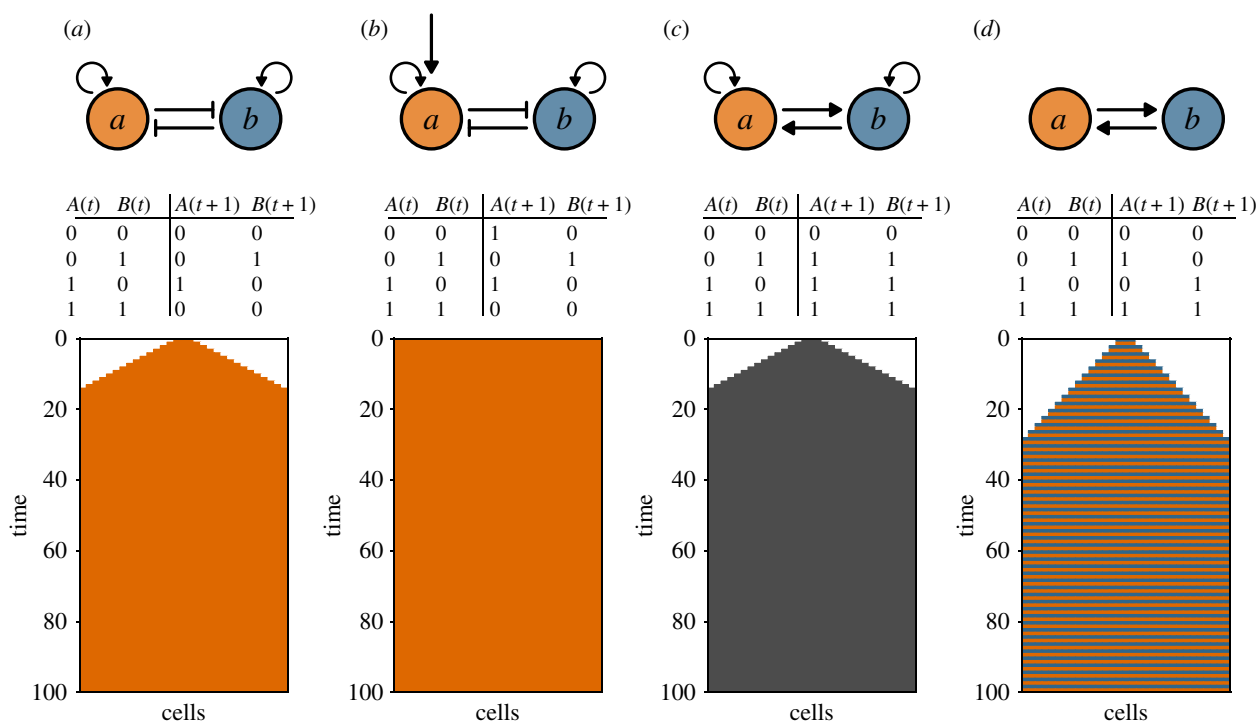


Figure 3. Some examples of patterns generated by two-gene Boolean GRNs. (a) A double-negative feedback loop, with self-activation, results in a pattern that expresses protein *A* (orange) stably, and expanding through the tissue. (b) A double-negative feedback loop with self-activation, where gene *a* is expressed constitutively, leads to the whole pattern expressing protein *A* stably through time. (c) A double-positive feedback loop with self-activation loops leads to both proteins *A* and *B* (grey) being expressed in the tissue in a stable way, and expanding through the tissue. (d) A double-positive feedback loop without self-activation leads to an alternating pattern where the tissue expresses first protein *A* (orange), then protein *B* (blue), and so on. Notice how the speed with which the pattern extends throughout the tissue is half the speed of patterns in *a* and *b*. This is because only protein *A* is allowed to propagate to the neighbouring cells (figure 2), so that the pattern can only extend when protein *A* is expressed.

evolution (including previous models of GRN evolution) would predict.

We performed Wright–Fisher evolutionary simulations with *toyLIFE* organisms in a strong selection, weak mutation regime (Methods), and selected the pattern in figure 3*b* as the evolutionary target—i.e. we assigned maximal fitness to it, and every other pattern became less fit as it differed more from the target (see Methods for the complete definition of the fitness function). We found that, after 100 000 mutations, 93% of simulations ended up finding one particular GRN among all 16 (GRN XI in electronic supplementary material, figure S10; see below), and the network in figure 3*b* (GRN V) only appears as the endpoint of evolutionary dynamics in three out of 10 000 simulations. In order to understand this somewhat unexpected result, we now discuss how Boolean GRNs are obtained from *toyLIFE* genotypes.

2.2. Regulation in *toyLIFE*

We will introduce gene regulation in *toyLIFE* through an example (for an in-depth discussion of *toyLIFE*'s rules, see Supplementary text in the electronic supplementary material and [26]). Consider the genotype in figure 4*a*. Proteins *A* and *B*, the expression products of genes *a* and *b*, respectively, bind together to form dimer *AB* (figure 4*b*). Because of *toyLIFE*'s interaction rules, the expression of gene *a* is activated by protein *A*, its own expression product. On the other hand, the expression of gene *b* is activated by the polymerase (it is a constitutively expressed gene), but it is inhibited by both proteins *A* and *B*. The dimer does not bind to any promoter (figure 4*c*). With this information, we can compute the

expression output of this genotype given each input, i.e. its truth table (figure 4*d*). When no protein is present, the polymerase (which is always present in the cell) will activate gene *b* and the output will consist of protein *B*. The same will happen if dimer *AB* is present in the cell: because it does not interact with either promoter, the polymerase will activate the expression of gene *b* again. If protein *A* is present, it will displace the polymerase and gene *b* will not be expressed, but *A* will also activate its own expression. Finally, if protein *B* is present, it will inhibit its own expression, and nothing will be expressed in the cell. In this way, we map a binary sequence (coding for the genome's two genes) into a Boolean GRN. It is interesting to note that this regulatory function cannot be expressed with an arrow diagram similar to those in figure 3: there is no way to represent the overriding effect that the dimer has on each protein's regulatory logic using this kind of diagram.

The cellular automaton is now uniquely determined by the GRN once we take into account two additional input states: protein *A* plus dimer, and protein *B* plus dimer. These two states can appear as a result of protein products propagating from one cell to the next. With this information we can unequivocally compute each genotype's corresponding cellular automaton.

It is worth noting that in the process of defining these phenotypic levels we have already introduced a lot of degeneracy. For instance, there are $2^{40} \approx 10^{12}$ genotypes with two genes, but they only give rise to 2152 different GRNs, which in turn generate only 453 different cellular automata—an average of $\approx 2 \times 10^9$ genotypes per cellular automaton. Not all GRNs are equally probable in genotype space, however: the

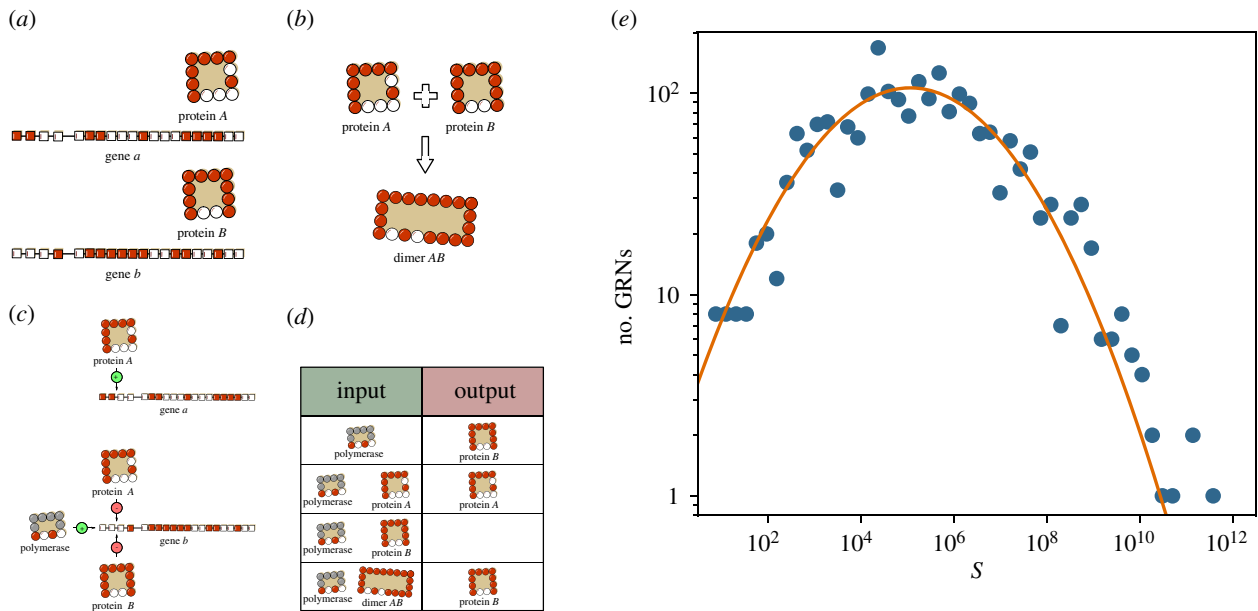


Figure 4. Regulatory logic in *toyLIFE*. (a) Example of a two-gene genotype in *toyLIFE*. These genes express strings of 16 amino acids that fold into a 4×4 lattice, following the rules of the HP model (Supplementary text in the electronic supplementary material). (b) Protein A and protein B can bind together to form dimer AB. (c) Regulatory logic of genes *a* and *b*. Protein A activates its own expression. The polymerase activates the expression of protein B, while both protein A and B inhibit it. The dimer AB does not bind either promoter. (d) Truth table representing the regulatory logic of this two-gene genotype, obtained from the information in (c). See text for details. (e) Not all GRNs generated by *toyLIFE* two-gene genotypes are equally likely in genotype space. In fact, the distribution of abundances (*S*) follows a log-normal distribution ($R^2 = 0.84$).

distribution of abundances of GRNs follows a log-normal distribution (figure 4e), which has been observed in many other genotype–phenotype models and has been shown to be universal under some very general assumptions [16,26–28]. The most abundant GRN is mapped by more than 500 billion genotypes, while the rarest one is mapped by only six genotypes. This phenomenon has been called phenotypic bias [15,16], and it is also observed in the distribution of abundances of cellular automata (electronic supplementary material, figure S11a). As a consequence, the 16 GRNs that generate the pattern in figure 3b (electronic supplementary material, figure S9) also have varying abundances. The most common one is GRN XI, mapped by 1.971×10^{11} genotypes, roughly 18% of all two-gene genotypes in *toyLIFE*. Its truth table appears in table 2. This is, admittedly, a very simple Boolean GRN, in which every input state leads to the same output: that of protein A being expressed. Previous work [29] has argued that simpler phenotypes should be more abundant in genotype space, and this is indeed what we observe at all phenotypic levels (electronic supplementary material, figure S12). In comparison, the least abundant Boolean GRN among these 16 (GRN VIII) is mapped by just 203 641 genotypes, a million times less abundant than GRN XI. Finally, the double-negative feedback loop in figure 3b that we were searching for originally (GRN V) is mapped by 9.4×10^6 genotypes, which is 5×10^{-5} times less abundant than GRN XI. As a result of this phenotypic bias in the 16 GRNs, when we evolve populations of *toyLIFE* organisms to express this simple pattern as described above, populations find GRN XI 93% of the times—although all 16 are equally fit in this scenario. In fact, the proportion of times our simulations end up in a particular GRN closely reflects its relative abundance in genotype space (electronic supplementary material, figure S13). In other words, introducing an additional level to the GRN-to-pattern genotype–phenotype map causes

Table 2. Truth table for GRN XI.

<i>A</i> (<i>t</i>)	<i>B</i> (<i>t</i>)	<i>A</i> (<i>t</i> + 1)	<i>B</i> (<i>t</i> + 1)
0	0	1	0
0	1	1	0
1	0	1	0
1	1	1	0

a bias in the abundances of different GRNs, which in turn affects evolutionary dynamics [30].

2.3. Pattern formation in *toyLIFE*

The differences in abundances in the Boolean GRNs generated by *toyLIFE* are magnified at the pattern level (electronic supplementary material, figure S10b): some patterns are mapped by billions of genotypes, while others are generated by only hundreds of them. This difference is critical, as we will see now. Suppose an evolutionary scenario where we select for pattern 113, as shown in figure 5a. This pattern is mapped by only 5312 genotypes, so finding it in genotype space seems hard *a priori*. However, naive evolutionary predictions would say that, being the fittest phenotype, it should be eventually selected and fixed in the population. When we perform the evolutionary simulations with pattern 113 as the target (Methods), it appears as the evolutionary endpoint in only 3% of the 1000 simulations. Instead, the pattern that appears in most of the simulations is pattern 109 (figure 5b), which has a fitness of 0.991 relative to that of pattern 113, and is mapped by 1.6×10^8 genotypes. There are 33 280 mutational paths between pattern 109 and pattern 113 (counted as the number of pairs of genotypes

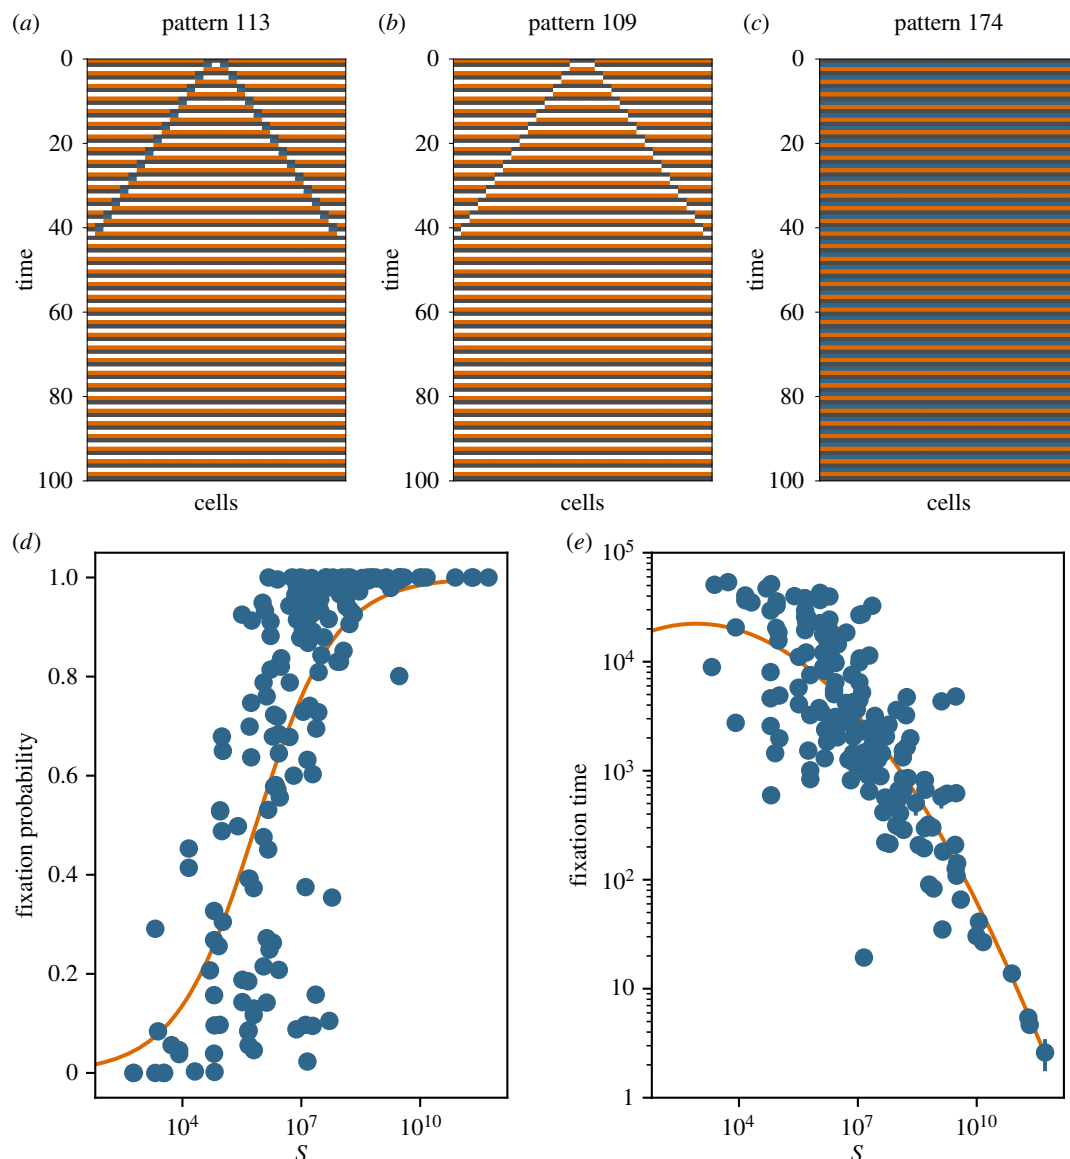


Figure 5. Evolving populations are not able to find rare patterns, even when they are fitter. (a) Pattern 113 is rarely found in our evolutionary simulations. (b) Pattern 109 is similar to 113, but it is generated by 1.64×10^8 genotypes—about 10^5 times commoner. As a result, it appears as the endpoint of our simulations 84% of the times. (c) Pattern 174 also appears as the endpoint of the simulations 8% of the time, even though it is not very similar to pattern 113. This is because of its high abundance in phenotype space: 1.36×10^8 genotypes are mapped to it. (d) This phenomenon is not restricted to pattern 113. The probability of finding a target pattern (p) goes to zero as the pattern abundance (S) decreases. Line: $p = (1 + (430767/S)^{1/2})^{-1}$, $R^2 = 0.58$. (e) Even when simulations do find the fittest pattern, the time to reach it (T) increases as pattern abundance decreases. Line: $\log_{10} T = 4.35 - 0.05(\log_{10} S - 2.93)^2$, $R^2 = 0.68$.

mapping to each pattern that are one point mutation apart), so populations expressing pattern 109 could eventually find pattern 113 without having to go through any fitness valley. However, this number represents only 0.0005% of all connections from pattern 109 to other phenotypes: finding pattern 113 from pattern 109 is truly like finding a needle in a haystack. In other words: the phenotypic bias towards pattern 109 is enough to counteract pattern 113's fitness benefit. Curiously enough, pattern 174 (figure 5c), which is not very similar to pattern 113, with a fitness of 0.54, also appears frequently as the endpoint of our simulations. In this case, there are no mutations from pattern 174 to 113, so it seems that some populations quickly find pattern 174 as a suboptimal fitness peak, and then become trapped in it, as there are no mutations to fitter alternatives.

This result means that some patterns will not be reachable by evolution, not because they are less fit, but because they are very rare in genotype space. This phenomenon is true for every

rare pattern, and indeed we see it in simulations where each of the 176 patterns obtained in our system is set as the target of evolution. The probability of finding the fittest pattern decreases dramatically with pattern abundance (figure 5a). And, even if the pattern is found, the time to find it decreases super-exponentially with pattern abundance (figure 5b).

3. Discussion and conclusion

The main intention of this work is to show that the complex mapping from DNA sequences to genetic circuits in real cells is in all likelihood biased towards some GRNs, so that some of them are much more common in genotype space. The results of our computational simulations show that this bias is enough to prevent populations from finding the fittest phenotype, consistent with previous results where bias affected evolutionary dynamics [15,31]. Several mechanisms had been previously

proposed to explain why populations do not reach the fittest solution, such as frequency-dependent selection [32] or the fittest versus the flattest [33,34]—which do not apply here, as our populations are always homogeneous. However, phenotypic bias is the first of these mechanisms that arises out of the intrinsic structure of the evolutionary search space, and it is completely independent from population effects and from the structure or function of the GRN being selected for. In this sense, phenotypic bias is playing the same role in evolutionary dynamics that entropy plays in statistical physics. The entropy of a macrostate is related to the number of microstates that are consistent with it without altering the properties that characterize the system. In statistical physics, macrostates are typically described by macroscopical properties such as temperature, pressure or volume, while microstates differ in the positions and velocities of individual particles. In evolutionary dynamics, there is a natural analogy between microstates and genotypes, on the one hand, and macrostates and phenotypes, on the other [35]. The conflict between energy and entropy found in physical systems is the same we have found between fitness and phenotypic bias, and the trapping by abundant phenotypes is akin to a glassy dynamics in physical systems [36].

Our results cannot be explained by phenotypic bias alone, however. In the simulations to find rare pattern 113, we found that a non-negligible fraction of simulations ended in pattern 174, which had a fitness of 0.54 but an abundance in genotype space that was similar to pattern 109, a fitter alternative. The reason populations got stuck in pattern 174 is because genotype space is structured as a complex network, and not all paths from one pattern to the other are actually possible. In this case, there are no connections between pattern 174 and either pattern 109 or 113, so, once the population has found this local fitness peak, there is no way it can reach the other, fitter alternatives under our selection regime. The effect of networked genotype spaces on evolutionary dynamics is far from trivial and has yet to be disentangled [37]. Further work has to be devoted to study its effects in this particular system.

The consequences of this work are immediate for the evolution of genetic circuits. Our results suggest that some ideal solutions could be hard to find in genotype space, and that evolution has had to work with more abundant, less efficient alternatives. However, the number of available phenotypes grows very quickly with genotype size in many computational genotype–phenotype maps [17,26,27], and so it is reasonable to expect that evolution could always find alternatives that are, if not optimal, at least highly functional. On the other hand, synthetic biologists trying to design a particular circuit could be aiming at a particularly rare structure, which would make its *a priori* evolution very unlikely. This would make that circuit very unstable in evolutionary terms, and mutations could easily change it into a different circuit, with undesired functions.

In relation to this, our results also suggest that phenotypic bias will have an effect on both robustness and evolvability. Previous models studying these properties in GRNs [11] found that they depended on the mechanism by which a GRN generates a pattern. Our results add a new layer, showing that more abundant GRNs will generate more robust patterns, independently of their mechanism or structure. Thus, understanding which GRNs are more abundant in genotype space is essential to unravel the evolution of robustness and evolvability.

We are aware of the limitations of toyLIFE as a discrete-time Boolean model to model continuous-time, stochastic protein concentration dynamics. However, phenotypic bias

is not a particular characteristic of toyLIFE and is rather very common in computational genotype–phenotype maps [15,16,29]. Thus, our main results are not limited to this particular choice of model, and they could easily be extended to other, more realistic genotype–phenotype maps. On the other hand, toyLIFE is a very convenient model to study multilevel genotype–phenotype relationships [26], which are complex and largely unknown. This model potential to generate complex behaviours is yet to be explored fully.

4. Methods and material

4.1. Fitness

The fitness function for our evolutionary simulations is calculated as follows: each pattern is a string in base 4 of length $L = 31 \cdot 100$. For every evolutionary scenario, we choose one particular pattern p_T as the target value, and assign fitness 1 to it. Then we compute the Hamming distance D of a pattern p to the target as

$$D(p, p_T) = \sum_{i=1}^L d_{p(i), p_T(i)}, \quad (4.1)$$

where $d_{i,j}$ is Kronecker's delta, which is equal to 1 if $i = j$ and 0 otherwise, and $p(i)$ is the i th letter in the string p . Fitness is then calculated as

$$f(p) = 1 - \frac{D(p, p_T)}{L}. \quad (4.2)$$

4.2. Evolutionary simulations

We assume a strong selection, weak mutation scenario. In this regime, Wright–Fisher dynamics are reduced to a continuous-time random walk in genotype space. We only consider point mutations, which arise in the population at constant rate μ , and the fixation rate of a new mutation is given by

$$\phi(f, N) = \mu N \frac{f - 1}{f^N - 1}, \quad (4.3)$$

where f is the fitness of the current phenotype relative to that of the mutant and N is population size [38]. We assume $\mu = 1$, which is equivalent to counting time in mutations instead of generations. Genotypes are binary strings of length 40, which are mapped to a pattern using toyLIFE 's rules (Supplementary text in the electronic supplementary material). We start the simulations choosing a genotype at random, and then simulate population dynamics using Gillespie's algorithm [39]. We simulated populations of size $N = 10\,000$ for $T = 100\,000$ mutations, and repeated this process for $R = 1000$ or $R = 10\,000$ replicates, depending on the experiment. The choice of population size was made so that deleterious mutations were hardly ever accepted.

4.3. Phenotypic complexity

For the results in electronic supplementary material, figure S12, we follow [29] and approximate the algorithmic complexity of a binary string $x = \{x_1, \dots, x_n\}$ as

$$\tilde{K}(x) = \begin{cases} \log_2(n) & x = 0^n \text{ or } 1^n \\ \log_2(n)^{\frac{1}{2}} [N_w(x_1, \dots, x_n) + N_w(x_n, \dots, x_1)] & \text{otherwise,} \end{cases} \quad (4.4)$$

where $n = |x|$ and $N_w(x)$ is the number of words in the dictionary created by the Lempel–Ziv algorithm [40]. For each phenotypic level (GRNs, cellular automata and patterns), we translate each base 4 string identifying the phenotype to binary code, and then compute \tilde{K} . So, for instance, string 312011 would become 110110000101. GRNs are represented as a binary string by reading all the output entries in the truth table (GRN II in electronic supplementary material, figure S10 is equivalent to 10001001) and then adding the output states of the two additional input states mentioned in the main text: protein A plus dimer, and protein B plus dimer. Therefore, GRNs can be uniquely represented as binary strings of length 12. Cellular automata are base 4 strings of length 64 that become binary strings of length 128 after converting from base 4 to base 2. Finally, patterns

are base 4 strings of length 3100 that become binary strings of length 6200.

Data accessibility. toyLIFE and all the code used to obtain the results in this paper are freely available at <https://github.com/pablocatalan/toylife/>.

Authors' contributions. P.C., S.M. and J.A.C. designed the work. P.C. carried out numerical simulations and wrote the first draft. All authors revised and worked on the final version of the manuscript.

Competing interests. We declare we have no competing interest.

Funding. P.C. is supported by a Ramón Areces Postdoctoral Fellowship. This research has been supported by Ministerio de Ciencia, Innovación y Universidades/FEDER (Spain/UE) through grant nos. PGC2018-098186-B-I00 (BASIC) and FIS2017-89773-P (MiMevo).

Acknowledgements. We are grateful to the suggestions of two anonymous referees that improved the manuscript's quality.

References

- Chen K, Rajewsky N. 2007 The evolution of gene regulation by transcription factors and microRNAs. *Nat. Rev. Genet.* **8**, 93–103. (doi:10.1038/nrg1990)
- Payne JL, Khalid F, Wagner A. 2018 RNA-mediated gene regulation is less evolvable than transcriptional regulation. *Proc. Natl Acad. Sci. USA* **115**, E3481–E3490. (doi:10.1073/pnas.1719138115)
- Davidson EH. 2010 *The regulatory genome: gene regulatory networks in development and evolution*. Burlington, MA: Academic Press.
- Alon U. 2006 *An introduction to systems biology: design principles of biological circuits*. London, UK: Chapman & Hall/CRC.
- Payne JL, Wagner A. 2015 Function does not follow form in gene regulatory circuits. *Sci. Rep.* **5**, 13015. (doi:10.1038/srep13015)
- Ahnert SE, Fink TMA. 2016 Form and function in gene regulatory networks: the structure of network motifs determines fundamental properties of their dynamical state space. *J. R. Soc. Interface* **13**, 20160179. (doi:10.1098/rsif.2016.0179)
- Santos-Moreno J, Schaerli Y. 2019 Using synthetic biology to engineer spatial patterns. *Adv. Biosyst.* **3**, 1800280. (doi:10.1002/adbi.201800280)
- Salazar-Ciudad I, Jernvall J, Newman SA. 2003 Mechanisms of pattern formation in development and evolution. *Development* **130**, 2027–2037. (doi:10.1242/dev.00425)
- Cotterell J, Sharpe J. 2010 An atlas of gene regulatory networks reveals multiple three-gene mechanisms for interpreting morphogen gradients. *Mol. Syst. Biol.* **6**, 425. (doi:10.1038/msb.2010.74)
- Schaerli Y, Munteanu A, Gili M, Cotterell J, Sharpe J, Isalan M. 2014 A unified design space of synthetic stripe-forming networks. *Nat. Commun.* **5**, 4905. (doi:10.1038/ncomms5905)
- Jiménez A, Cotterell J, Munteanu A, Sharpe J. 2015 Dynamics of gene circuits shapes evolvability. *Proc. Natl Acad. Sci. USA* **112**, 2103–2108. (doi:10.1073/pnas.1411065112)
- Jiménez A, Cotterell J, Munteanu A, Sharpe J. 2017 A spectrum of modularity in multi-functional gene circuits. *Mol. Syst. Biol.* **13**, 925. (doi:10.15252/msb.20167347)
- Banzhaf W, Kuo PD. 2004 Network motifs in natural and artificial transcriptional regulatory networks. *J. Biol. Phys. Chem.* **4**, 85–92. (doi:10.4024/2040405.jbpc.04.02)
- Khatri BS, McLeish TC, Sear RP. 2009 Statistical mechanics of convergent evolution in spatial patterning. *Proc. Natl Acad. Sci. USA* **106**, 9564–9569. (doi:10.1073/pnas.0812260106)
- Schaper S, Louis AA. 2014 The arrival of the frequent: how bias in genotype-phenotype maps can steer populations to local optima. *PLoS ONE* **9**, e86635. (doi:10.1371/journal.pone.0086635)
- Dingle K, Schaper S, Louis AA. 2015 The structure of the genotype–phenotype map strongly constrains the evolution of non-coding RNA. *Interface Focus* **5**, 20150053. (doi:10.1098/rsfs.2015.0053)
- Louis AA. 2016 Contingency, convergence and hyper-astronomical numbers in biological evolution. *Stud. Hist. Philos. Sci. C* **58**, 107–116.
- Kauffman SA. 1969 Metabolic stability and epigenesis in randomly constructed genetic nets. *J. Theor. Biol.* **22**, 437–467. (doi:10.1016/0022-5193(69)90015-0)
- Payne JL, Moore JH, Wagner A. 2014 Robustness, evolvability, and the logic of genetic regulation. *Artif. Life* **20**, 111–126. (doi:10.1162/ARTL_a_00099)
- Ingalls BP. 2013 *Mathematical modeling in systems biology: an introduction*. Cambridge, MA: MIT Press.
- Bornholdt S. 2008 Boolean network models of cellular regulation: prospects and limitations. *J. R. Soc. Interface* **5**, S85–S94. (doi:10.1098/rsif.2008.0132.focus)
- Wagner A. 2011 *The origins of evolutionary innovations*. Oxford, London: Oxford University Press.
- Istrail S, Davidson EH. 2005 Logic functions of the genomic cis-regulatory code. *Proc. Natl Acad. Sci. USA* **102**, 4954–4959. (doi:10.1073/pnas.0409624102)
- Wolfram S. 2002 *A new kind of science*, vol. 5. Champaign, IL: Wolfram Media.
- Arias CF, Catalán P, Manrubia S, Cuesta JA. 2014 toyLIFE: a computational framework to study the multi-level organisation of the genotype–phenotype map. *Sci. Rep.* **4**, 7549. (doi:10.1038/srep07549)
- Catalán P, Wagner A, Manrubia S, Cuesta JA. 2018 Adding levels of complexity enhances both robustness and evolvability in a multi-level genotype–phenotype map. *J. R. Soc. Interface* **15**, 20170516. (doi:10.1098/rsif.2017.0516)
- Manrubia S, Cuesta JA. 2017 Distribution of genotype network sizes in sequence-to-structure genotype–phenotype maps. *J. R. Soc. Interface* **14**, 20160976. (doi:10.1098/rsif.2016.0976)
- García-Martín JA, Catalán P, Manrubia S, Cuesta JA. 2018 Statistical theory of phenotype abundance distributions: a test through exact enumeration of genotype spaces. *Europhys. Lett.* **123**, 28001. (doi:10.1209/0295-5075/123/28001)
- Dingle K, Camargo CQ, Louis AA. 2018 Input–output maps are strongly biased towards simple outputs. *Nat. Commun.* **9**, 761. (doi:10.1038/s41467-018-03101-6)
- Greenbury SF, Schaper S, Ahnert SE, Louis AA. 2016 Genetic correlations greatly increase mutational robustness and can both reduce and enhance evolvability. *PLoS Comput. Biol.* **12**, e1004773. (doi:10.1371/journal.pcbi.1004773)
- Johnston IG, Ahnert SE, Doye JP, Louis AA. 2011 Evolutionary dynamics in a simple model of self-assembly. *Phys. Rev. E* **83**, 066105. (doi:10.1103/PhysRevE.83.066105)
- Ayala FJ, Campbell CA. 1974 Frequency-dependent selection. *Annu. Rev. Ecol. Syst.* **5**, 115–138. (doi:10.1146/annurev.es.05.110174.000555)
- Wilke CO, Wang JL, Ofria C, Lenski RE, Adami C. 2001 Evolution of digital organisms at high mutation rates leads to survival of the flattest. *Nature* **412**, 331–333. (doi:10.1038/35085569)

34. Beardmore RE, Gudelj I, Lipson DA, Hurst LD. 2011 Metabolic trade-offs and the maintenance of the fittest and the flattest. *Nature* **472**, 342–346. (doi:10.1038/nature09905)
35. Sella G, Hirsh AE. 2005 The application of statistical physics to evolutionary biology. *Proc. Natl Acad. Sci. USA* **102**, 9541–9546. (doi:10.1073/pnas.0501865102)
36. Götze W. 2008 *Complex dynamics of glass-forming liquids: a mode-coupling theory*, vol. 143. Oxford, UK: Oxford University Press.
37. Aguirre J, Catalán P, Cuesta JA, Manrubia S. 2018 On the networked architecture of genotype spaces and its critical effects on molecular evolution. (<http://arxiv.org/abs/1804.06835>)
38. Ewens WJ. 2004 *Mathematical population genetics*. New York, NY: Springer.
39. Gillespie DT. 1976 A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J. Comput. Phys.* **22**, 403–434. (doi:10.1016/0021-9991(76)90041-3)
40. Lempel A, Ziv J. 1976 On the complexity of finite sequences. *IEEE Trans. Inf. Theory* **22**, 75–81. (doi:10.1109/TIT.1976.1055501)