

Robustness and stability of spin-glass ground states to perturbed interactions

Vaibhav Mohanty^{1,2,*} and Ard A. Louis^{1,†}

¹*Rudolf Peierls Centre for Theoretical Physics, University of Oxford, Oxford, OX1 3NP, United Kingdom*

²*MD-PhD Program and Program in Health Sciences and Technology, Harvard Medical School, Boston, Massachusetts 02125, USA and Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*



(Received 10 December 2020; revised 20 October 2022; accepted 16 December 2022; published 18 January 2023)

Across many problems in science and engineering, it is important to consider how much the output of a given system changes due to perturbations of the input. Here, we investigate the glassy phase of $\pm J$ spin glasses at zero temperature by calculating the robustness of the ground states to flips in the sign of single interactions. For random graphs and the Sherrington-Kirkpatrick model, we find relatively large sets of bond configurations that generate the same ground state. These sets can themselves be analyzed as subgraphs of the interaction domain, and we compute many of their topological properties. In particular, we find that the robustness, equivalent to the average degree, of these subgraphs is much higher than one would expect from a random model. Most notably, it scales in the same logarithmic way with the size of the subgraph as has been found in genotype-phenotype maps for RNA secondary structure folding, protein quaternary structure, gene regulatory networks, as well as for models for genetic programming. The similarity between these disparate systems suggests that this scaling may have a more universal origin.

DOI: [10.1103/PhysRevE.107.014126](https://doi.org/10.1103/PhysRevE.107.014126)

I. INTRODUCTION

Systems in which the input can be represented as a sequence of characters appear across science and engineering and especially commonly in biology and computer science. In these fields, notions of system robustness or sensitivity can be defined to quantify the outputs' average resistance to small changes in the input sequences.

In biology, the mapping from genotypes (which store the information) to phenotypes (which describe biological properties) can be abstracted as genotype-phenotype maps (GP maps). Examples include four-letter RNA sequences and 20-letter protein sequences that can be mapped to their physical folded states, and gene-regulatory networks, which can, for example, be described by Boolean networks [1] where a set of weights represent the gene interaction strengths. The responses of such systems to changes in the input sequences have been extensively studied computationally and analytically [2–22].

For GP maps, an important concept is the set of genotypes (sequences) that map to a particular phenotype, often called the neutral set. These present a number of commonalities across GP maps [12,13,18]: The neutral sets are typically highly connected so that they can be viewed as networks, which can be traversed by single mutational steps, leading to enhanced evolvability [4,19]. Neutral sets are also called neutral networks. In some cases such as RNA GP maps, the neutral network can split into smaller component networks that are disconnected due to biophysical constraints [5,20].

The neutral set size (NSS) can vary over many orders of magnitude, and is typically strongly biased, with a small fraction of the phenotypes taking up the majority of genotypes. Such phenotype bias can strongly affect evolutionary outcomes [8,11,21,23].

A key property of the neutral set for GP maps is the mutational robustness, typically defined in the literature as the fraction ρ_p of single-character mutations in a genotype that produce the same phenotype p , averaged over the neutral set of all genotypes that produce phenotype p . That is to say, for a biological input-output map $f(g)$, which takes in a genotype g of d characters chosen from an alphabet $K = \{K_0, \dots, K_{k-1}\}$, the robustness of phenotype p is defined as

$$\rho_p(f) = \frac{1}{|\mathcal{G}_p|d(k-1)} \sum_{g \in \mathcal{G}_p} n_{p,g}, \quad (1)$$

where \mathcal{G}_p is the neutral set of all genotypes g whose output is the phenotype p , and $n_{p,g}$ is the number of nearest neighbors of g mapping to p defined as the number of genotypes g' satisfying $f(g') = f(g) = p$ that differ from g by a Hamming distance of 1. Thus, $\rho_p(f) \in [0, 1]$ measures the mean probability that a mutation from $g \in \mathcal{G}_p$ to a neighboring genotype $g' \in \mathcal{G}_p$ results in the same phenotype p .

What all the biological models mentioned above hold in common is that they each have high levels of mean robustness to changes in their inputs, excepting pathological or adversarial examples. To quantify this statement, consider the naive uncorrelated expectation that $\rho_p \approx |\mathcal{G}_p|/k^d$ because, in a large randomly assigned GP map, the probability that a nearest-neighboring genotype yields phenotype p is approximately equal to the probability that any genotype drawn at

*mohanty@hms.harvard.edu

†ard.louis@physics.ox.ac.uk

random from the entire input space yields phenotype p . This probability is $|\mathcal{G}_p|/k^d$, so $n_{p,g} \approx d(k-1)|\mathcal{G}_p|/k^d$, and $\rho_p \approx |\mathcal{G}_p|/k^d$. And indeed this is true for completely uncorrelated and randomly assigned GP maps [12]. However, biology- and computer science-inspired GP maps such as the RNA secondary structure maps, protein folding maps, gene regulatory networks, linear genetic programs, and digital logic gate maps exhibit a remarkably similar observed scaling for robustness $\rho_p \propto \log |\mathcal{G}_p|$ rather than $\rho_p \propto |\mathcal{G}_p|$ [18] over a range of $|\mathcal{G}_p|$, which can span several orders of magnitude in biological GP maps [12]. Such enhanced mutational robustness is necessary for the neutral sets to percolate [12], which, in turn, is critical for the evolutionary process because it allows for neutral exploration of the neutral set, greatly enhancing the ability of a population to find novel phenotypic variation [4]. This enhanced robustness also facilitates the exploration of fitness landscapes [24].

Robustness has a direct counterpart in computer science, the sensitivity, which measures the likelihood that flipping a single input bit will alter the output bit of Boolean functions $f: \{0, 1\}^d \rightarrow \{0, 1\}$ that map binary sequences of length d onto a single binary output. In other words, low sensitivity corresponds to high robustness. Huang's [25] famously short solution to the decades-old sensitivity conjecture [26] concerning induced subgraphs of the n -dimensional hypercube graph has recently brought a great deal of attention to the sensitivity analysis of Boolean functions. While scaling laws for sensitivity are typically not measured in the manner done for biological robustness, there are many mathematical equivalencies between the two. In the literature (see, e.g., Refs. [25–28]), quantities are defined such as local sensitivity, function sensitivity, and average sensitivity, which are similar to biological robustness. In this paper, we do not explicitly use these definitions from sensitivity analysis, instead opting for the GP map-inspired ones. But, we point out that there is a seemingly understudied connection between robustness in biological systems and sensitivity in Boolean functions.

Given the wide range of systems for which high robustness is observed, we ask here whether a similar phenomenon can be found in spin glasses, which have a rich history in statistical and condensed matter physics. They have been intensely studied since the 1970s [29,30], and have led to many important insights in physics and other related disciplines, including computer science [31–33]. More recently, the spin glass Hamiltonian has been used as a phenomenological model for epistatic genotype-to-fitness landscapes in which different sites (e.g., DNA, genes expressions, or amino acids) may couple to each other [34–40]. An important application has been to viruses [35–39]. By taking sequence data over time, inverse statistical physics methods can be employed to learn the interactions (bond configurations) between different sites. In the context of evolution, therefore, one can interpret the ground state of the spin glass energy landscape as the global fitness peak on an evolutionary fitness landscape. The interactions between spins in such a system can depend on a number of biological or environmental factors [34,38]. The interesting question in the context of fitness landscapes is, what is the robustness of the sequences of spins to mutations in the bond configurations (the epistatic couplings)?

In this paper, we consider this question by investigating the spin glass phase at $T = 0$ by calculating the robustness of $\pm J$ spin glass ground-state configurations to a sign flip perturbation of a single bond. Our investigation is related to the concept of spin glass bond chaos or disorder chaos [41], which considers how the ground state of a spin glass changes when all the couplings are perturbed by a small amount; notable investigations (see Refs. [41–45]) have been conducted on continuous J spin glasses, whereas we consider the $\pm J$ spin glasses. Another similar concept is the recently introduced σ criticality [46], which considers the effects on the ground state of only a single-bond perturbation, as we do here, but again this is a continuous J spin glass. Our use of $\pm J$ spin glasses is an important distinction between our work and these previous studies because the discreteness of the bond configuration domain of the $\pm J$ spin glass, modeled as a hypercube graph, allows us to investigate the topological features of the subgraphs formed by bond configurations mapping to the same ground state. The bond chaos and σ -criticality investigations are mainly interested in how the spin glass ground states change, and we consider these aspects here as well, but in this work we are mainly interested in the universality of the subgraph network features in the bond configuration domain.

In addition to spin glasses on random graphs, we examine special cases, namely the Sherrington-Kirkpatrick model and one-dimensional (1D) Edwards-Anderson model. Various network topological properties are computed for interaction domain subgraphs comprised of bond configurations that all map to the same ground-state spin configuration. We find that these subgraphs obey the same logarithmic scaling law between robustness and neutral set size as the analogous biological and computer science GP maps described above, suggesting that this high robustness may hold for a much wider set of physical systems.

II. MODEL AND METHODS

A. Spin glass model

Consider an undirected, unweighted random graph $G(V, E)$ with vertex set V and edge set E . We place Ising spins on each vertex, and each edge represents a nonzero interaction between spins. A spin configuration $s \in \{\pm 1\}^{|V|}$ can be written as a sequence of $+1$ and -1 values, so it is essentially a binary sequence of length $|V|$. A set of interactions (bond configuration) $J \in \{\pm 1\}^{|E|}$ similarly is a sequence of $+1$ and -1 values of length $|E|$. The spin glass Hamiltonian

$$\mathcal{H}_G(s; J) = - \sum_{\{i,j\} \in E} J_{ij} s_i s_j - \sum_{i \in V} h_i s_i \quad (2)$$

contains couplings between all spins, which are connected by an edge in G . The single-spin, external magnetic field interactions h_i are independently and identically distributed uniformly on the interval $[-10^{-4}, +10^{-4}]$, noting that the external field $h_i \ll J_{ij} \in \{-1, +1\}$; we incorporate an external field in order to break the possible degeneracies of the spin glass ground state. In our numerical simulations described below, we find that this choice of distribution for h_i is sufficient to ensure a unique ground state for all of our simulations.

The input-output map considered in our study is the spin glass ground-state optimization function

$$\Gamma_G : \{\pm 1\}^{|E|} \rightarrow \{\pm 1\}^{|V|} \quad (3)$$

defined for the graph G . For a bond configuration J , $\Gamma_G(J)$ outputs the ground-state configuration s that minimizes the Hamiltonian Eq. (2). The most common task in spin glass theory is to find s given a particular bond configuration J . In this paper, we study an inverse problem, namely the relationship between the set of all bond configurations $\{J\}$ that generate a particular output s .

To efficiently represent this system, we note that the collection of all binary sequences of length n can be represented by a n -dimensional undirected hypercube graph $Q_n(U, F)$. This is accomplished by mapping each binary sequence to a vertex in $Q_n(U, F)$ and placing edges between two vertices if the corresponding sequences have a Hamming distance of 1 between them. The hypercube graph has vertex set U with $|U| = 2^n$ and edge set F with $|F| = 2^{n-1}n$.

The domain of Γ_G accordingly has a mapping to the $|E|$ -dimensional hypercube graph $Q_{|E|}(U, F)$. In general, for graphs G that can produce geometrical frustration in the spin glass, $\Gamma_G(J)$ follows no pattern and is difficult to calculate [47], even more so because of the degeneracy-breaking external random field interactions $\{h_i\}$. But, because of frustration, two sets of bond configurations $J^{(i)}$ and $J^{(j)}$ corresponding to adjacent vertices in $Q_n(U, F)$ often have $\Gamma_G(J^{(i)}) = \Gamma_G(J^{(j)})$. The vertices corresponding to all J such that $\Gamma_G(J) = s$ for some fixed s induce a subgraph $H_s(U_s, F_s)$ of $Q_{|E|}$. It follows that $\bigcup_s U_s = U$. In the GP map literature, a neutral network (or neutral set) for phenotype (here, the ground state) s is the graph $H_s(U_s, F_s)$.

In this paper, we numerically compute topological properties of each neutral network of spin glass bond configurations which all map to the same ground-state spin configuration. We consider multiple random graphs for G as well as the fully connected graph; the latter case is known as the Sherrington-Kirkpatrick (SK) model of a spin glass [29]. In our simulations, we also impose the condition that every node in G has at least one neighbor. Additionally, in this paper we also consider the one-dimensional Edwards-Anderson (EA) model [30], for which the relationship between robustness $|F_s|$ and vertex count $|U_s|$ (equivalent to robustness) becomes analytically solvable. Our simulated spin glasses are relatively small because (i) we need to calculate the ground states exactly for each bond configuration, (ii) every possible bond configuration is considered for each spin glass in order to accurately determine neutral network properties. In other words, the number of times Eq. (2) is computed in order to find all ground states for a random graph $G(V, E)$ scales as $O(2^{|V|} \times 2^{|E|})$, which forces us to use relatively small systems. Our largest model ($|V| = 8$, $|E| = 28$) involves calculation of an exact ground state for every single-bond configuration (over 2.6×10^8 bond configurations); this is a similar size to systems for which exact ground-state calculations like ours were carried out in the context of the spin glass literature (see, e.g., Ref. [48]).

B. Definitions of topological quantities

The following parameters are computed for the neutral networks in bond configuration space:

Robustness. The neighbor count $n_{p,g}$ in Eq. (1) is equivalent to the degree of a vertex $g \in U_s$ within a neutral network $H_s(U_s, F_s)$. The mean degree is related to the number of edges by

$$\sum_{v \in U_s} \deg(v) = 2|F_s|. \quad (4)$$

Here, we compute the robustness ρ_s , simply dividing the above quantity by the size of the subgraph $|U_s|$ and by the length of the input sequence $|E|$:

$$\rho_s = \phi_{ss} \equiv \frac{2|F_s|}{|U_s||E|} \in [0, 1]. \quad (5)$$

The notation ϕ_{ss} will become clear below when we also treat the transition probability ϕ_{rs} of a bond perturbation leading to a different ground state r . It is important to notice that the robustness is nothing other than the edge-to-vertex ratio $|F_s|/|U_s|$ normalized to the range $[0, 1]$ and also is equivalently the normalized mean degree within the neutral network H_s . In many real-world GP maps including RNA and protein folding, Boolean threshold networks, and genetic algorithms, it has been observed that $\rho_s \sim \log |U_s|$ or equivalently $|F_s| \sim |U_s| \log |U_s|$. We will test this scaling for spin glass systems.

Transition probability. We define the transition probability ϕ_{rs} to be the average probability that, given a bond configuration which maps to ground state s , a single-bond perturbation changes the ground state from s to r . Graph theoretically, we can think of this in the following way: consider two neutral networks of $Q_{|E|}(U, F)$ called $H_s(U_s, F_s)$ and $H_r(U_r, F_r)$ so that s and r are two different spin configurations ($s \neq r$). Let $T_{rs} = T_{sr}$ be the set of edges connecting H_s and H_r . The transition probability of $s \rightarrow r$ due to a single-bond perturbation is

$$\phi_{rs} \equiv \frac{|T_{rs}|}{|U_s||E|} \in [0, 1], \quad s \neq r. \quad (6)$$

In many real-world GP maps, it has been observed that often $\phi_{rs} \sim |U_r|$, or equivalently $|T_{rs}| \sim |U_r||U_s|$ [8,12]. We will also test this scaling for spin glass systems.

We now see that, for the case where $s = r$, the transition probability T_{ss} , based on the definition above, simply counts the number of edges within the induced subgraph $H_s(U_s, F_s)$, and therefore $T_{ss} = F_s$. In the definition of ϕ_{ss} in Eq. (5), an additional factor of 2 comes from the double counting of edges; recall, the robustness is equivalent to the (normalized) mean degree, which means that the number of edges must be double counted. From the definitions, it is easy to check that the normalization condition $\sum_r \phi_{rs} = 1$ holds, where the sum includes $r = s$.

Rank-size plot. A rank-size plot is useful for visualizing how many orders of magnitude the neutral set sizes span. It may be used to deduce if there is a power-law (i.e., generalized Zipf's law) relationship between rank and neutral set size, as is seen for some GP maps.

Degree distribution, clustering coefficients, assortativity, and betweenness centrality. We also compute the degree distribution, clustering coefficients, assortativity, and betweenness




FIG. 1. $G(V, E)$ for (a) and (b) two representative examples of random graphs and (c) and (d) two representative examples of complete graphs (Sherrington-Kirkpatrick model). Each Ising spin s_i is placed on a vertex, and each interaction J_{ij} is placed on an edge. Each spin also experiences an external random field h_i .

centrality for the neutral networks studied here. Degree distributions give us information about the modality and skew of the degrees of the vertices in the neutral networks. Clustering coefficients, which turn out to be trivially zero for these systems, are informative about the number of neighboring nodes, which share a neighbor. Assortativity measures correlation between the degrees of neighboring vertices. Betweenness centrality of a node measures the number of shortest paths in the network passing through that node; it is often compared with, and is expected to be positively correlated with, the degree of the node [49]. Because our systems are small, generalizable trends regarding these quantities are more difficult to resolve, but calculating them assists our physical intuition, so we present them in the Supplemental Material [50] along with their mathematical definitions.

III. RESULTS

We now present results for the properties described in the previous section for the neutral networks $H_s(U_s, F_s)$ for spin glasses defined on random graphs as well as for the SK model. Two instances of random graphs ($|V| = 7, |E| = 15$ and $|V| = 9, |E| = 15$) and two SK models of differing sizes ($|V| = 6, |E| = 15$ and $|V| = 8, |E| = 28$) are used as the representative simulation examples for this main text; these graphs G are shown in Fig. 1. To demonstrate the consistency of our results across many instances of the random graphs, properties of additional random graph instances and their network topological properties are described in the plots in the Supplemental Material [50]. We also present the 1D EA model as a special case where we can calculate the exact relationship between robustness and neutral network size.

A. Robustness

For all topologies for G , we find that each induced subgraph $H_s(U_s, F_s)$ has exactly one connected component, regardless of size. Whether this will also hold for larger systems is unclear, but this lack of multiple component networks is different from, for example, the RNA GP map, where biophysical constraints (mainly that $GC \leftrightarrow UA$ and $CG \leftrightarrow AU$ bond changes are not possible by single-point mutations) lead to fragmentation of the neutral networks [5,20]. In Fig. 2, we plot the robustness of neutral networks versus the neutral set size.

Each plot point represents a particular ground state and its associated neutral network, and the plot axes have been set to ensure that all neutral networks appear in the viewing window. Due to the presence of the random fields $\{h_i\}$, one expects to observe at most $2^{|V|-1}$ unique ground states obtained from all bond configurations. Our numerical simulations find in most cases that indeed $2^{|V|-1}$ are found, but in a small number of $|V| = 9$ simulations we find half of this number, namely $2^{|V|-2}$ ground states, possibly due to some symmetry of graph G which is not necessarily easily apparent from inspection (see Supplemental Material [50] for an example).

The spin glasses on random and complete graphs all show behavior consistent with the $\rho_s \propto \log |U_s|$ relationship, which is also seen for the closely related scaling of robustness with neutral set size found for many GP maps. For these small models, the range of variation in $|U_s|$ seems to depend on the density of bonds in graph G . For the $|V| = 7, |E| = 15$ random graph, neutral set sizes ranged from 216–1200 bond configurations ($0.0066 \leq |U_s|/2^{|E|} \leq 0.037$), for the $|V| = 9, |E| = 15$ random graph, neutral set sizes ranged from 64–928 bond configurations ($0.0020 \leq |U_s|/2^{|E|} \leq 0.028$), for the $|V| = 6, |E| = 15$ SK model, neutral set sizes ranged from 608–1616 bond configurations ($0.019 \leq |U_s|/2^{|E|} \leq 0.049$), and for the $|V| = 8, |E| = 28$ SK model, neutral set sizes ranged from 1,118,512–3,754,464 bond configurations ($0.0044 \leq |U_s|/2^{|E|} \leq 0.014$). As the density increases, with the SK model being the limiting case for a fixed $|V|$, the range of $|U_s|$ spans less than half an order of magnitude, while the less densely connected $|V| = 9, |E| = 15$ random graphs can have over an order of magnitude of variation in $|U_s|$, reaching nearly two orders of magnitude in variation (see Supplemental Material [50] for an example). This is notably smaller than the range of variation found in biological GP maps, where maps with a similar number of genotypes can have more orders of magnitude of variation. For instance, the RNA secondary structure GP map with length 12 has $4^{12} \approx 1.7 \times 10^7$ genotypes and a neutral set size variation of nearly five orders of magnitude when including the nonfolding phenotype [5]. In comparison, our $|V| = 8$ SK model has $2^{28} \approx 2.7 \times 10^8$ genotypes (bond configurations) but shows less than half an order of magnitude of variation in $|U_s|$.

While the range of variation in $|U_s|$ for these small spin glass systems is too small to fully confirm the expected log

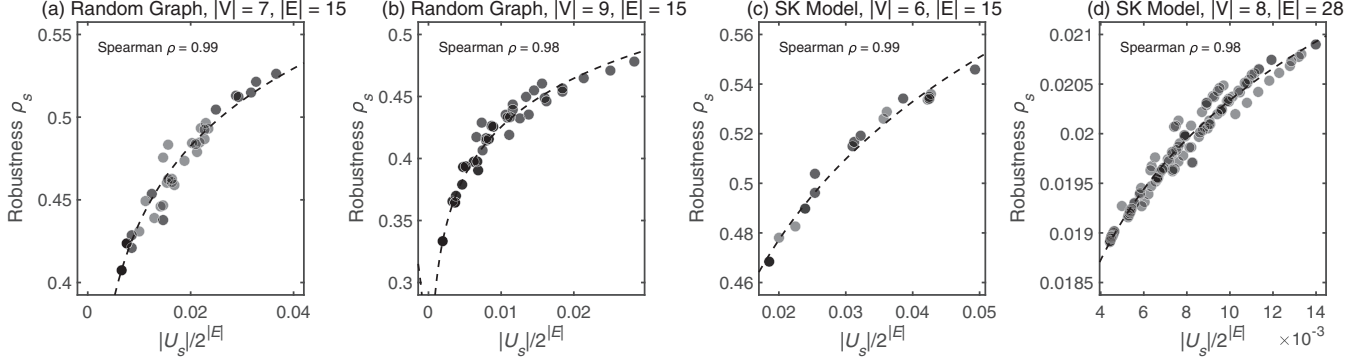


FIG. 2. Linear-linear plot of robustness ρ_s versus normalized neutral set size $|U_s|/2^{|E|}$ for representative random graphs (a) and (b) and representative SK models (c) and (d). The dashed curve arises from the linear least-squares fit calculated for data on a linear-log scale (see also Fig. 3). The data are consistent with a logarithmic relationship between robustness and NSS. The plot window is adjusted to ensure that plot points for all neutral networks are within the viewing window.

scaling, the data in Fig. 3 support the notion that the robustness is still enhanced when compared to an uncorrelated random null model. This null model is calculated by taking each spin glass ground-state mapping $\Gamma_G(J)$ and randomizing the input-output pairings while keeping the subgraph size the same. A single-bond perturbation $J \mapsto J'$ will result in a spin configuration $s = \Gamma_G(J')$ being selected with probability $|U_s|/2^{|E|}$, regardless of $\Gamma_G(J)$. Thus, $\phi_{rs} \approx |U_r|/2^{|E|}$, even for $r = s$, in the null model.

B. Transition probabilities

Transition probabilities are plotted in Fig. 4 for the largest neutral network of each spin glass model. As found for biologically inspired GP maps [8,12], ϕ_{rs} is typically much smaller than the ρ_s ; this is apparent from comparing the vertical axes in Figs. 4 and 3. As a null model, we again use

$$\phi_{rs} \approx |U_r|/2^{|E|}, \quad r \neq s. \quad (7)$$

Overall, for the random graphs, as can be seen in Fig. 4, this null model curve matches the linear least-squares fit, suggesting that vertices of subgraphs $H_r(U_r, F_r)$ (for $r \neq s$) with nonzero transition probability are approximately randomly

distributed with frequency $\approx |U_r|/2^{|E|}$ in the neighborhoods of all vertices $v \in |U_s|$. To interpret these findings about robustness and transition probabilities, consider a vertex $v \in H_s(U_s, F_s)$. In its set of nearest neighbors, this vertex is expected to see an overrepresentation (relative to the null model) of other vertices belonging to $H_s(U_s, F_s)$, and it sees a random assortment of other vertices $u \in H_r(U_r, F_r)$ for various $r \neq s$ with probability proportional to $|U_r|/2^{|E|}$. These findings are in accordance with other GP map studies.

For the SK models, there is less agreement between the null expectation and the observed lines of best fit. We speculate that the increased bond density in G is responsible for this effect on transition probabilities, but our systems are currently too small to investigate this effect.

C. Size-rank distributions

In the GP map literature, there has been a lot of interest in phenotype bias, the observation that the neutral set sizes can vary over many orders of magnitude, which can determine evolutionary outcomes [11,21] even when natural selection is also at play. In Fig. 5 we show that such phenotype bias also exists for spin glass systems. As we already mentioned

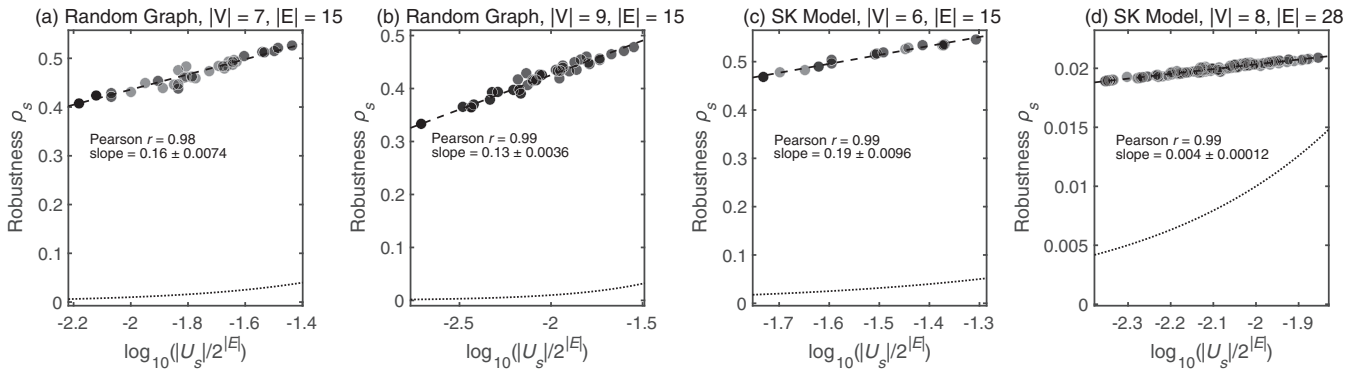


FIG. 3. Linear-log plot of robustness ρ_s versus normalized neutral set size $|U_s|/2^{|E|}$ for representative random graphs (a) and (b) and representative SK models (c) and (d). The dashed line is the line of best of fit in the linear-log scale. The data are consistent with a linear relationship between robustness and the log of NSS. Regression means and 95% confidence intervals are given for the slope of the best fit lines. The dotted line shows the random null expectation for robustness. The plot window is adjusted to ensure that plot points for all neutral networks are within the viewing window.




FIG. 4. Transition probability ϕ_{rs} (for $r \neq s$ versus normalized neutral set size $|U_s|/2^{|E|}$) for the largest neutral network s of representative random graphs (a) and (b) and representative SK models (c) and (d). The transition probability ϕ_{rs} is, starting with a bond configuration that maps to ground state s , the probability that a single bond perturbation would now cause a change to ground state r , with $r \neq s$. To see the $r \neq s$ transition probabilities compared to the $r = s$ case, see the Supplemental Material [50]. Here, the dashed line is the line of best of fit in the linear-log scale; the dotted line is the identity line (in which abscissa = ordinate). In (d), for the SK model, the line of best fit deviates most from the identity line (which is outside the viewing area), but the data are consistent with a linear fit, as for the other models. Regression means and 95% confidence intervals are given for the slope of the best fit lines.

in the robustness section, the range of variation in $|U_s|$ for our small spin glass systems is observed to be smaller than what the biological GP maps exhibit. Nonetheless, the rank plots show a consistent behavior independent of spin glass graph topology G . Recent work on GP maps has suggested that there are two main classes of rank plots [9,14,51]. For the first class, the distribution of neutral set sizes obeys a Zipf-like power law, which arises from models in which input site ordering is strongly constrained (including Boolean neural networks [52,53]). The other class is a log-normal distribution, which appears, for example, in RNA secondary structure GP maps [11,51]. An open question is whether or not these spin glass systems also fall into one of these two classes. The current systems are still too small to conclusively answer this question, but the log-log rank versus size plots in Supplemental Material [50] suggest a deviation from Zipf's law.

D. Effects of bond perturbation on ground-state spin configuration

Our single-bond perturbation numerical experiments are informative regarding the nature of the ground state in the $T = 0$ glassy phase. Our robustness calculations have already

shown that a single-bond perturbation often leads to no change in the ground state, much more often than what the random null model predicts. Given that the robustness is equivalent to the normalized mean degree within a neutral network, we can understand robustness at the individual bond level by looking at the relationship between the degree of an individual vertex within the neutral network and various physical parameters such as ground-state energy. We can ask, how is the degree within a neutral network related to the ground-state energy? How many spin flips occur in the ground-state energy when a single bond is perturbed? Is the unperturbed ground state a local minimum in the energy landscape generated by the new, perturbed bond configuration?

From the degree distributions within neutral networks with $|E| = 15$ (shown in the Supplemental Material [50]), we see that the range of degrees for the vertices in most neutral networks tends to span from 4 up to 15. In the Supplemental Material [50], we empirically confirm the expected result that the degree of a vertex is positively correlated with its betweenness centrality; that is, there tends to be a larger number of shortest paths passing through vertices with higher degree. Intuitively, one should then expect that vertices with higher degree have ground states, which are more energetically stable




FIG. 5. Plot of normalized induced subgraph size $\log_{10}(|U_s|/2^{|E|})$ (equivalent to the NSS or frequency) versus the rank of the size for (a) and (b) random graphs and (c) and (d) SK models. All the models exhibit a similar rapid decay of the neutral set size.

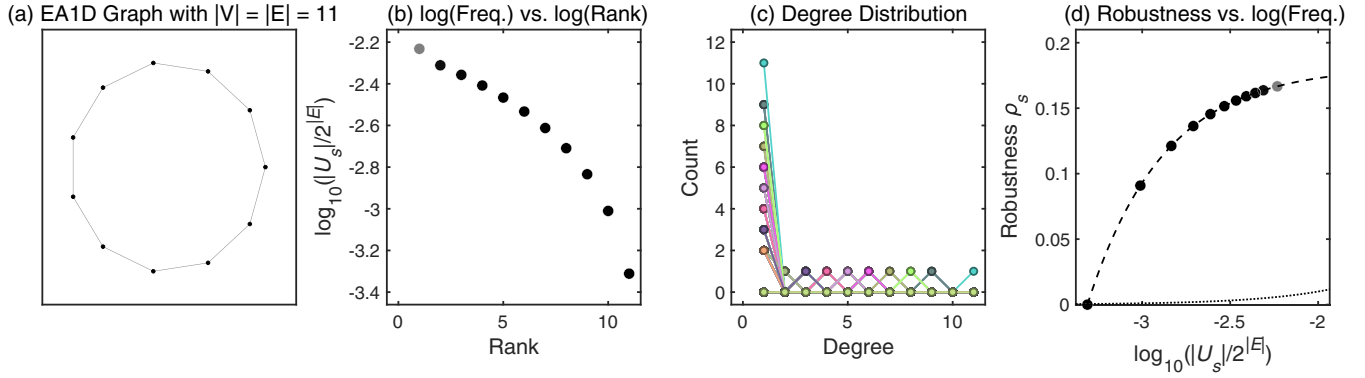


FIG. 6. Results for 1D Edwards-Anderson model: (a) Graph representation of the 1D Edwards-Anderson model with $|V| = 11$, (b) neutral set size versus rank plot on log-log scale, (c) degree distribution of neutral set vertices, and (d) robustness ρ_s versus log of the normalized neutral set size $|U_s|/2^{|E|}$. The dashed line is the analytical result from Eq. (8).

to bond perturbation. In other words, the ground-state energy is sufficiently low, and the minimum is deep such that a single-bond perturbation tends not to change the ground state. This trend is what we see in the numerical results presented in the Supplemental Material [50], where we also discuss the different scenarios when ground-state change due to bond perturbation does occur.

E. Analytically tractable special case: 1D Edwards-Anderson model

The Edwards-Anderson (EA) model is a special case, which deserves individual treatment. The EA model, the original theory of spin glasses [30], is simply a spin glass on a lattice with nearest-neighbor interactions only.

For the 1D Edwards-Anderson (EA) model with periodic boundary conditions, the topology of which is shown in Fig. 6(a), the behavior of ρ_s is analytically tractable. Let us call a bond configuration with an even number of antiferromagnetic interactions ($J_{ij} = -1$) an even bond configuration and a bond configuration with an odd number of antiferromagnetic interactions an odd bond configuration.

Every even bond configuration has an exactly determinable ground state with no spin frustration, which can be found using the following algorithm:

1. Choose an arbitrary site i , and set the i th spin to $s_i = +1$ or $s_i = -1$ arbitrarily.
2. Determine s_{i+1} by setting $s_{i+1} = J_{i,i+1}s_i$, so $s_{i+1} = s_i$ if $J_{i,i+1} = +1$ or $s_{i+1} = -s_i$ if $J_{i,i+1} = -1$.
3. Continue to s_{i+2} by similarly setting $s_{i+2} = J_{i+1,i+2}s_{i+1}$, and continue around the entire spin chain until we have set s_{i-1} . At this point one can check and confirm that $-J_{i-1,i}s_{i-1}s_i = -1$, so all bonds are satisfied, and each pairwise interaction decreases the energy by the same amount ($= -1$).
4. If $\sum_{i=1}^{|V|} h_i s_i < 0$, then all spins need to be flipped $s_i \rightarrow -s_i$. If $\sum_{i=1}^{|V|} h_i s_i > 0$, then all the spins are kept as is. The ground-state energy is exactly $E_{gs} = -|V| - \sum_{i=1}^{|V|} h_i s_i$, and we have now ensured $\sum_{i=1}^{|V|} h_i s_i > 0$.

Having an even number of antiferromagnetic interactions will ensure that the last spin s_{i-1} does not experience frustration, and the bond between the last spin s_{i-1} and the first spin s_i

contributes a negative amount to the energy: $-J_{i-1,i}s_{i-1}s_i = -1$. On the other hand, odd bond configurations will always have one bond at which there is frustration, and there will be a positive contribution to the energy $-J_{i-1,i}s_{i-1}s_i = +1$. The ground state for odd bond configurations will be identical to the ground state of one of the neighboring even bond configurations, but the energy will, of course, be higher. Exactly which neighboring bond configuration that is depends on the random fields h_i .

Because an odd bond configuration has the same ground state as one of its neighboring even bond configurations, the neutral network $H_s(U_s, F_s)$ for every ground state must be a star graph $S_{|U_s|-1}$ that has $|U_s|$ nodes, with one central node corresponding to an even bond configuration, and $|U_s| - 1$ peripheral nodes corresponding to odd bond configurations. The number of nodes of each star graph neutral network depends on the random fields h_i . Figure 6(c) shows that the degree distribution matches our theory that all neutral networks are star graphs; for each neutral network, there is one node with $|U_s| - 1$ neighbors, and $|U_s| - 1$ nodes with one neighbor.

To calculate robustness, we note that the star graph $S_{|U_s|-1}$ has $|U_s|$ nodes and $|F_s| = |U_s| - 1$ edges. Plugging $|F_s|$ into Eq. (5) immediately gives us that the robustness is

$$\rho_s = \frac{1}{|U_s||E|} \sum_{v \in U_s} \deg(v) = \frac{2}{|E|} \left(1 - \frac{1}{|U_s|} \right). \quad (8)$$

The points in Fig. 6(d) fall exactly along this curve. The maximum number of nodes in the star graph is $|U_s| = |E| + 1 = |V| + 1$ because each even bond configuration has at most $|E| = |V|$ neighboring odd bond configurations. A size-rank plot is also shown for the 1D EA model in Fig. 6(b).

The 1D EA model serves as an example in which linear-log scaling is not seen for ρ_s . Nevertheless, the relationship between robustness and subgraph size found in the 1D EA model exhibits a ρ_s that is significantly higher than would be expected from a random mapping Γ_G of inputs to outputs, so this 1D model still exhibits enhanced robustness. For the 2D EA model we could not find an analytically tractable ρ_s .

IV. DISCUSSION

Our main result is that, by studying the mapping from spin glass bond configurations to the $T = 0$ ground state, we observe several properties that are also observed for other input-output maps, such as the GP maps found in biology and computer science. These include (i) redundancy, in that many genotypes in the context of GP maps, or bond configurations in this case, map the same output (a phenotype for the GP maps, or a ground state for the spin glasses), (ii) phenotype frequency bias, in that the number of bond configurations (genotypes) mapping to the ground state (phenotype) vary significantly, and (iii) enhanced mutational robustness, in that the probability of a single-bond perturbation changing the ground state is larger than expected from a random null model, and moreover that the robustness of a ground state to single-bond perturbations often scales in the same way with the logarithm of the NSS (the number of bond configurations mapping to a ground state) as found for biological GP maps [9,10,12,16,17]. Another interesting result, also seen for GP maps [8,12] is that, in contrast to the robustness, the transition probabilities, defined as the likelihood of a flip of the spin yielding a different ground state, do scale proportionally to the NSS, as one would expect from a random model. The similarity to the GP map behavior suggests that there may be a more universal argument (based, for example, on algorithmic information theory [54,55]) for these scaling properties.

Our spin glass models are relatively small because finding the ground state of a spin glass is typically computationally expensive and scales badly with system size. Depending on graph topology, finding the ground state of a spin glass can be NP hard [47]. Knowledge that the robustness of a ground state is large may potentially offer improvements to ground-state-finding algorithms by providing a measure of stability of certain ground states as a function of parameter space. It would also be interesting to check some of our results on

significantly larger graphs. We find, for example, that all our subgraphs that map to the same ground state form only a single component; that is, they are connected by single-bond perturbations. Will this percolation property hold for larger systems, or will these subgraphs start to fragment?

Mapping epistatic interactions onto spin glass Hamiltonians to derive sequence-to-fitness maps has been especially important for the study of viral evolution [35–39]. These models typically have continuous J , so it will be interesting to see if the kind of scaling properties of robustness that we find here for the $\pm J$ spin glasses carry over to these more complex systems. If, as we expect, a concept akin to high robustness persists, then this may have implications for the stability of fitness peaks to environmental changes, and may affect how easy it is to find continuously increasing paths to fitness maxima [24].

Another potentially interesting future direction of research is to explore these results about robustness in the complementary setting of the sensitivity of Boolean functions [25–28]. It would be interesting to see whether similar high robustness-low sensitivity results can be found in this related context.

ACKNOWLEDGMENTS

The authors would like to thank Wilber Lim, Mehrana Nejad, Sloan Nietert, and Shuofeng Zhang for helpful discussions, and two anonymous referees for careful review and many helpful insights. V.M. has been supported by a Marshall Scholarship and by Awards No. T32GM007753 and No. T32GM144273 from the National Institute of General Medical Sciences. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute of General Medical Sciences, the National Institutes of Health, or the Marshall Aid Commemoration Commission.

-
- [1] S. Kauffman, Homeostasis and differentiation in random genetic control networks, *Nature (London)* **224**, 177 (1969).
 - [2] A. Wagner, Distributed robustness versus redundancy as causes of mutational robustness, *BioEssays* **27**, 176 (2005).
 - [3] A. Wagner, *Robustness and Evolvability in Living Systems*, 3rd ed., Princeton studies in complexity (Princeton University Press, Princeton, 2007).
 - [4] A. Wagner, Robustness and evolvability: A paradox resolved, *Proc. R. Soc. B* **275**, 91 (2008).
 - [5] J. Aguirre, J. M. Buldú, M. Stich, and S. C. Manrubia, Topological structure of the space of phenotypes: The case of RNA neutral networks, *PLoS ONE* **6**, e26324 (2011).
 - [6] J. L. Payne and A. Wagner, Constraint and contingency in multifunctional gene regulatory circuits, *PLoS Comput. Biol.* **9**, e1003071 (2013).
 - [7] J. L. Payne, J. H. Moore, and A. Wagner, Robustness, evolvability, and the logic of genetic regulation, *Artificial Life* **20**, 111 (2014).
 - [8] S. Schaper and A. A. Louis, The arrival of the frequent: How bias in genotype-phenotype maps can steer populations to local optima, *PLoS ONE* **9**, e86635 (2014).
 - [9] S. F. Greenbury, I. G. Johnston, A. A. Louis, and S. E. Ahnert, A tractable genotype-phenotype map modelling the self-assembly of protein quaternary structure, *J. R. Soc. Interface* **11**, 20140249 (2014).
 - [10] S. F. Greenbury and S. E. Ahnert, The organization of biological sequences into constrained and unconstrained parts determines fundamental properties of genotype-phenotype maps, *J. R. Soc. Interface* **12**, 20150724 (2015).
 - [11] K. Dingle, S. Schaper, and A. A. Louis, The structure of the genotype-phenotype map strongly constrains the evolution of non-coding RNA, *Interface Focus* **5**, 20150053 (2015).
 - [12] S. F. Greenbury, S. Schaper, S. E. Ahnert, and A. A. Louis, Genetic correlations greatly increase mutational robustness and can both reduce and enhance evolvability, *PLoS Comput. Biol.* **12**, e1004773 (2016).

- [13] S. E. Ahnert, Structural properties of genotype–phenotype maps, *J. R. Soc. Interface* **14**, 20170275 (2017).
- [14] M. Weiß and S. E. Ahnert, Phenotypes can be robust and evolvable if mutations have non-local effects on sequence constraints, *J. R. Soc. Interface* **15**, 20170618 (2018).
- [15] D. Nichol, M. Robertson-Tessi, A. R. A. Anderson, and P. Jeavons, Model genotype–phenotype mappings and the algorithmic structure of evolution, *J. R. Soc. Interface* **16**, 20190332 (2019).
- [16] C. Q. Camargo and A. A. Louis, Boolean threshold networks as models of genotype-phenotype maps, *Complex Networks XI*, 143 (2020).
- [17] T. Hu, M. Tomassini, and W. Banzhaf, A network perspective on genotype–phenotype mapping in genetic programming, *Genetic Program. Evol. Mach.* **15**, 375 (2020).
- [18] S. Manrubia, J. A. Cuesta, J. Aguirre, S. E. Ahnert, L. Altenberg, A. V. Cano, P. Catalán, R. Diaz-Uriarte, S. F. Elena, J. A. García-Martín, P. Hogeweg, B. S. Khatri, J. Krug, A. A. Louis, N. S. Martin, J. L. Payne, M. J. Tarnowski, and M. Weiß, From genotypes to organisms: State-of-the-art and perspectives of a cornerstone in evolutionary dynamics, *Phys. Life Rev.* **38**, 55 (2021).
- [19] J. L. Payne and A. Wagner, The causes of evolvability and their evolution, *Nature Rev. Genet.* **20**, 24 (2019).
- [20] S. Schaper, I. G. Johnston, and A. A. Louis, Epistasis can lead to fragmented neutral spaces and contingency in evolution, *Proc. R. Soc. B* **279**, 1777 (2012).
- [21] K. Dingle, F. Ghaddar, P. Šulc, and A. A. Louis, Phenotype bias determines how natural rna structures occupy the morphospace of all possible shapes, *Mol. Biol. Evol.* **39**, msab280 (2022).
- [22] A. H. Wright and C. L. Laue, Evolving complexity is hard, [arXiv:2209.13013](https://arxiv.org/abs/2209.13013).
- [23] I. G. Johnston, K. Dingle, S. F. Greenbury, C. Q. Camargo, J. P. K. Doye, S. E. Ahnert, and A. A. Louis, Symmetry and simplicity spontaneously emerge from the algorithmic nature of evolution, *Proc. Natl. Acad. Sci. USA* **119**, e2113883119 (2022).
- [24] S. F. Greenbury, A. A. Louis, and S. E. Ahnert, The structure of genotype-phenotype maps makes fitness landscapes navigable, *Nat Ecol. Evol.* **6**, 1742 (2022).
- [25] H. Huang, Induced subgraphs of hypercubes and a proof of the Sensitivity Conjecture, *Ann. Math.* **190**, 949 (2019).
- [26] N. Nisan and M. Szegedy, On the degree of boolean functions as real polynomials, *Comput. Complexity* **4**, 301 (1994).
- [27] A. Bernasconi, Sensitivity vs. block sensitivity (an average-case study), *Inf. Process. Lett.* **59**, 151 (1996).
- [28] S. Chakraborty, On the sensitivity of cyclically-invariant Boolean functions, *Discrete Math. Theor. Comput. Sci.* **13**, 51 (2011).
- [29] D. Sherrington and S. Kirkpatrick, Solvable Model of a Spin-Glass, *Phys. Rev. Lett.* **35**, 1792 (1975).
- [30] S. F. Edwards and P. W. Anderson, Theory of spin glasses, *J. Phys. F* **5**, 965 (1975).
- [31] M. Mezard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond*, World Scientific lecture notes in physics Vol. 9 (World Scientific, Singapore, 1987).
- [32] M. Mezard and A. Montanari, *Information, Physics, and Computation*, Oxford graduate texts (Oxford University Press, Oxford, 2009).
- [33] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing: An Introduction*, International series of monographs on physics No. 111 (Oxford University Press, Oxford, 2001).
- [34] Y. Guo, M. Vucelja, and A. Amir, Stochastic tunneling across fitness valleys can give rise to a logarithmic long-term fitness trajectory, *Sci. Adv.* **5**, eaav3842 (2019).
- [35] R. H. Y. Louie, K. J. Kaczorowski, J. P. Barton, A. K. Chakraborty, and M. R. McKay, Fitness landscape of the human immunodeficiency virus envelope protein that is targeted by antibodies, *Proc. Natl. Acad. Sci. USA* **115**, E564 (2018).
- [36] T. C. Butler, J. P. Barton, M. Kardar, and A. K. Chakraborty, Identification of drug resistance mutations in HIV from constraints on natural evolution, *Phys. Rev. E* **93**, 022412 (2016).
- [37] J. P. Barton, N. Goonetilleke, T. C. Butler, B. D. Walker, A. J. McMichael, and A. K. Chakraborty, Relative rate and location of intra-host HIV evolution to evade cellular immunity are predictable, *Nature Commun.* **7**, 11660 (2016).
- [38] K. Shekhar, C. F. Ruberman, A. L. Ferguson, J. P. Barton, M. Kardar, and A. K. Chakraborty, Spin models inferred from patient-derived viral sequence data faithfully describe HIV fitness landscapes, *Phys. Rev. E* **88**, 062705 (2013).
- [39] T. A. Hopf, J. B. Ingraham, F. J. Poelwijk, C. P. I. Schärfe, M. Springer, C. Sander, and D. S. Marks, Mutation effects predicted from sequence co-variation, *Nature Biotechnol.* **35**, 128 (2017).
- [40] S. Cocco, C. Feinauer, M. Figliuzzi, R. Monasson, and M. Weigt, Inverse statistical physics of protein sequences: A key issues review, *Rep. Prog. Phys.* **81**, 032601 (2018).
- [41] A. J. Bray and M. A. Moore, Chaotic Nature of the Spin-Glass Phase, *Phys. Rev. Lett.* **58**, 57 (1987).
- [42] T. Aspelmeyer, Free-Energy Fluctuations and Chaos in the Sherrington-Kirkpatrick Model, *Phys. Rev. Lett.* **100**, 117205 (2008).
- [43] M. Ney-Nifle and A. P. Young, Chaos in a two-dimensional Ising spin glass, *J. Phys. A: Math. Gen.* **30**, 5311 (1997).
- [44] M. Sasaki, K. Hukushima, H. Yoshino, and H. Takayama, Temperature Chaos and Bond Chaos in Edwards-Anderson Ising Spin Glasses: Domain-Wall Free-Energy Measurements, *Phys. Rev. Lett.* **95**, 267203 (2005).
- [45] W. Wang, J. Machta, and H. G. Katzgraber, Bond chaos in spin glasses revealed through thermal boundary conditions, *Phys. Rev. B* **93**, 224414 (2016).
- [46] C. M. Newman and D. L. Stein, Ground-state stability and the nature of the spin glass phase, *Phys. Rev. E* **105**, 044132 (2022).
- [47] F. Barahona, On the computational complexity of Ising spin glass models, *J. Phys. A: Math. Gen.* **15**, 3241 (1982).
- [48] S. Boettcher and T. M. Kott, Exact enumeration of ground states in the Sherrington-Kirkpatrick spin glass, *Phys. Rev. B* **72**, 212408 (2005).
- [49] S. Oldham, B. Fulcher, L. Parkes, A. Arnatkeviciute, C. Suo, and A. Forn, Consistency and differences between centrality measures across distinct classes of networks, *PLOS ONE* **14**, e0220061 (2019).
- [50] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevE.107.014126> for extended data and discussion.
- [51] S. Manrubia and J. A. Cuesta, Distribution of genotype network sizes in sequence-to-structure genotype–phenotype maps, *J. R. Soc. Interface* **14**, 20160976 (2017).

- [52] C. Van den Broeck and R. Kawai, Learning in feedforward Boolean networks, *Phys. Rev. A* **42**, 6210 (1990).
- [53] G. Valle-Pérez, C. Q. Camargo, and A. A. Louis, Deep learning generalizes because the parameter-function map is biased towards simple functions, *7th International Conference on Learning Representations* (New Orleans, 2019).
- [54] K. Dingle, C. Q. Camargo, and A. A. Louis, Input–output maps are strongly biased towards simple outputs, *Nature Commun.* **9**, 761 (2018).
- [55] K. Dingle, G. V. Pérez, and A. A. Louis, Generic predictions of output probability based on complexities of inputs and outputs, *Sci. Rep.* **10**, 4415 (2020).