

Dynamic Landscapes: A Model of Context and Contingency in Evolution

Laura Feeney,¹ David V. Foster,² Jacob G. Foster,³ Tanja Gesell,⁴ Antony Millner,⁵ and Molly Rorick⁶

¹*Swedish Institute of Computer Science and Uppsala University, Sweden*

²*Institute for Biocomplexity and Bioinformatics, University of Calgary, Canada*

³*Complexity Science Group, Department of Physics and Astronomy, University of Calgary, Canada*

⁴*Center for Integrative Bioinformatics Vienna, Max F. Perutz Laboratories, University of Vienna, Austria*

⁵*Centre for the Environment, University of Oxford, UK*

⁶*Department of Genetics, Yale University, USA*

(Dated: August 15, 2008, update September 4, 2008)

The basic mechanics of evolution have been understood since Darwin. But debate continues over the relative importance of the structure of genotype space and of ecological interaction in driving macroevolutionary phenomena. In this paper we propose a simple, abstract model capturing some key features of fitness-landscape and ecological models of evolution. Our model describes evolutionary dynamics in high-dimensional, structured genotype space with a significant role for interspecific interaction — thus incorporating context and contingency. We find some promising qualitative similarity with the empirical facts about macroevolution, including broadly distributed extinction sizes and realistic exploration of the genotype space. The abstraction of our model permits numerous interpretations and applications beyond macroevolution, including molecular evolution and technological innovation.

I. INTRODUCTION

What drives evolution? Since the days of Darwin, the prevailing explanation has emphasized heritable variation and selection. But while the mechanism of heredity and the importance of random mutation for generating variation have both been thoroughly explicated, the nature and causal agents of selection remain rather mysterious. As a result, we still struggle to explain those most striking events in the drama of life: the mass extinctions, adaptive radiations, and even the prosaic local speciations that generate the millions or tens of millions of extant species (which are, nevertheless, but a hundredth of the number that at one time or another occupied this planet) [1].

Theories tackling these (macro)evolutionary questions fall into two broad categories, distinguished both in their originating impulse and their central metaphor. The first approach, coming from molecular genetics, focuses on the fitness landscape introduced by Sewall Wright [2]. For selection cannot act on genotype directly; rather, genotypes are selected based on the phenotype they generate. The fitness landscape metaphor therefore focuses on the genotype-phenotype/fitness map, that is, the assignment of some scalar fitness value to every point in the genotype space. This fitness value in turn influences the population-genetic process of mutation and selection occurring on the genotype space. Mutations to higher fitness genotypes are favored, with lower fitness mutants tending to extinction with high probability. The second approach, coming from ecology and dynamical systems theory, emphasizes the role of interactions in driving species to extinction. Here the genotype space is essentially ignored, and the fate of species in some (often fixed) ecology is determined by interacting population dynamics (e.g. the replicator dynamics of evolutionary game theory [3]), with extinct species usually replaced

by a random extant species e.g. [4, 5]. As an example, abundant prey may drive the growth of a predator population, which in turn drives some third prey species to extinction.

Both approaches capture important qualitative features of macroevolution. Given recent insight into the subtleties of the genotype-phenotype map, e.g. its many-to-one character, it seems clear that the interaction between genotype (the realm of heredity and variation) and phenotype (the realm of selection) is of central importance to phenomena like speciation, adaptive radiation (e.g. the Cambrian explosion [6]), and punctuation in the rate of evolution [4, 7]. Likewise, variants of the ecological approach have established an important role for endogenous interspecies dynamics in generating extinction, up to and including mass extinctions that eliminate most biodiversity. In other words, both models capture important aspects of context (e.g. particular ecological configurations rendering some genotypes unfit) and contingency (e.g. particular mutational histories limiting or enhancing accessible genotypes) in evolution. In this paper we marry the stylized facts captured by the two approaches into a simple, unified model of evolutionary dynamics. Briefly, our model describes evolution on a richly structured fitness landscape, where context and contingency strongly determine the subsequent evolution of the simulated biosphere. Our model captures some qualitative features of empirical evolution, including broad distributions of extinction sizes and the rather surprising result that evolution takes place as an advancing front through genotype space.

In Section II we summarize the Fitness Landscape and Ecological Models discussed in the Introduction. In particular, we gather the stylized facts that guide the construction of our model. In Section III we describe the model and the numerical methods used in our simulations. Section IV reviews the results, and in Section V

we discuss these results and offer our conclusions as well as an outline of future work.

II. BACKGROUND

Our model of evolutionary dynamics lies at the nexus of two important traditions in the mathematical modeling of evolution. In both cases, researchers have sought to distill invariant and universal features of the evolutionary process into mathematics. In reviewing such models, we similarly seek a list of essential “stylized facts” which can inform the construction of a parsimonious, realistic abstraction of the evolutionary process.

A. Fitness Landscapes

Since their introduction by Sewall Wright [2], fitness or adaptive landscapes have played a dominant role in evolutionary theory [1]. This dominance follows from their exceeding conceptual simplicity: the genetic code of the organism defines a space of potential genetic configurations or genotype space (denoted \mathcal{G} henceforth) on which is defined a fitness function Φ (related to the likelihood of survival). This fitness function is a map $\Phi : \mathcal{G} \rightarrow \mathbb{R}^+$ from genotype to some scalar measure of fitness. A population of individuals is then defined over \mathcal{G} , the population dynamics of which are influenced by $\Phi(\mathcal{G})$.

In much of the literature following Wright’s original suggestion, the “landscape” is literally modeled on a landscape: these so called “rugged landscapes” have many adaptive peaks, of various heights, separated by adaptive valleys [1]. But this simplification of Wright’s picture has several major flaws. First, in low-dimensional models (which, as suggested by direct interpretation of the landscape metaphor, look like real-world landscapes), it is rather difficult to generate speciation with any reasonable probability. Since selection pushes the population up adaptive peaks and away from adaptive valleys, crossing a fitness valley to a new (perhaps better) peak is unlikely. Shifting balance, wherein the population is subdivided so that stochastic shift across a valley is more likely and higher fitness types can then sweep the population, fails to explain the observed fecundity of the biosphere [1]. Founder effect speciation, wherein a small number of individuals found a new, small population in which the crossing probability is enhanced, is likewise unsuccessful [1].

The second flaw also follows directly from the landscape picture. For the low-dimensional picture suggested by real peaks and valleys is grossly inaccurate. \mathcal{G} is in fact enormously high dimensional, as most organisms have thousands of genes and millions or billions of base-pairs (hence $\dim \mathcal{G} \sim 10^6 - 10^9$) [1]. Fisher already observed that this high dimensionality should convert “adaptive peaks” into saddle points, and favored a single peaked landscape, albeit in enormously high dimensions [8].

Third, Kimura’s claim that most evolutionary change is neutral (i.e. indifferent with respect to fitness [9]) seems to have been at least partially validated by extensive experimental evidence of neutrality. For example, the genotype/phenotype map for RNA and proteins is now known to be many-to-one, implying that many sequences are selectively neutral (here the phenotype seen by selection is the fold of the RNA or protein) [7]. Neutrality if taken to its extreme would lead to a totally flat fitness landscape rather than a rugged one; here selection would play no role whatsoever.

An important compromise embracing ruggedness, high dimensions, and neutrality was suggested by John Maynard Smith: “It follows that if evolution by natural selection is to occur, functional proteins must form a continuous network which can be traversed by unit mutational steps without passing through nonfunctional intermediaries” [10]. The essence of this suggestion — that genotype space is percolated by a network or networks of more-or-less equally fit genotypes, which nevertheless represent a small fraction of all possible genotypes — forms the core of the neutral network or holey landscape approach pioneered theoretically by Gavrillets [1, 8, 11, 12]. Selection plays a role, defining the neutral network(s) and keeping populations from mutating into the “holes” of the landscape (or alternately driving rapid evolution out of the holes onto the ridges). But so does neutrality, as most evolution takes place neutrally along the networks (which through their interweaving define the adjacency and accessibility of various protein folds or morphological types associated to the networks). Some advocates of the neutral network picture go so far as to claim that the structuring of the genotype space by these neutral networks plays a primary role in shaping the phenomena of speciation, adaptive radiation, and punctuated equilibrium [7]. Even adopting a moderate version of this view makes clear that the continuous, topologically trivial fitness landscapes of the “peaks and valleys” picture are a far cry from the actual complexity of accessible genotype space.

The holey landscape picture is amply supported by both theoretical evidence (in which neutral networks seem to be an inevitable consequence of a surprising variety of model specifications) and by empirical evidence from studies of RNA and proteins [1, 11, 12]. Later in the paper we shall present an extended interpretation of our model in terms of both proteins and RNA, in which specific evidence is reviewed. Theoretically, we follow the very simplest variant of the holey landscape, which Gavrillets calls Russian roulette: each genotype is assigned a fitness of 1 with probability p and 0 with probability $1 - p$ [1]. Summarizing the stylized facts: we want

- selection to matter, while ignoring small differences in fitness;
- very few genotypes to be fit;
- the genotype space to be suitably high dimensional;

- neutrality to play a substantial role;
- neutral *networks* to exist in the genotype space.

In Section III it will become clear that these facts lead naturally to several components of our model.

B. Ecological Models

While fitness landscape models focus on genotype space and the way it is structured by selection (hence highlighting the puzzle of speciation), so-called ecological models focus on interspecific interaction (hence highlighting the puzzle of extinction, both background and mass [4, 5]). Evidence for extinction played an important role in overturning the static pre-Darwinian biology. More recently, the discovery of mass extinctions (most famously the KT extinction that wiped out the dinosaurs) raised questions of causation that were originally answered exogenously. For example, there is ample evidence for the asteroid impact now thought to have precipitated the KT extinction (if not caused it entirely) [5]. Physics models of self-organized criticality suggested that these mass extinctions could also be caused by the same mechanism that operates at small scales, i.e. ecological interaction between species [4, 5].

The Bak-Sneppen model is perhaps the most famous “ecological model” [13]. Here N species are arranged in a circle, with each species i assigned a “fitness” B_i , which represents a barrier to mutation. Because of exponential separation of timescales under a barrier-dependent probability of mutation $p_i \sim e^{-B_i/T}$ (with T setting the timescale of mutations) the species with the lowest barrier B_{low} at some timestep is always assumed to mutate. This species and its two neighbors are assigned new, random fitness values (the mutation in B_{low} presumably alters the fitness landscape of its neighbors). The model self-organizes into a critical state characterized by scale-invariant power laws, with the exponents of the power laws sensitive to the dimension of the lattice defining the interaction structure [4]. For example, if the lifetime of a species is defined as the time between two mutations at that site, then the distribution of lifetimes t is $N(t) \sim t^{-\alpha}$ with $\alpha = 1.1$ for the 1-d lattice (the empirical distribution for fossil genera has $\alpha = 2$) [4]. Now define an avalanche as a series of mutations that are causally related (i.e. a change in species i ’s barrier makes j ’s the least fit, causing it to change, in turn causing one of *its* neighbors to change, etc.). The size of an avalanche s is thus the number of causally related mutations forming a given avalanche. The distribution of avalanche sizes also obeys a power law $N(s) \sim s^{-\tau}$ with $\tau = 1.1$ [4] (the empirical $\tau \approx 2$ [5]).

Solé, Manrubia and collaborators defined a somewhat more realistic ecological model of extinction and speciation in which there is some notion of inheritance as well as population dynamics [5]. In this model they define a connectivity matrix γ_{ij} valued on the interval $(-1, 1)$.

These matrix elements represent the interactions between species, and are not strictly speaking “food web” interactions but rather some generalized positive or negative influence of the species j on the species i . Species viability is binary: $S_i = 0$ or $S_i = 1$ (extinct or extant, respectively). Using a step function $\Phi(z) = 1, z > 0$, they define dynamics on this ecology [5]:

$$S_i(t+1) = \Phi \left(\sum_{j=1}^n \gamma_{ij}(t) S_j(t) \right) \quad (1)$$

The model proceeds by randomly varying γ_{ij} (external driving); implementing the population dynamics, which may render some species extinct; and replacing extinct species with ϵ -varied mutants of an extant species randomly chosen to undergo adaptive radiation. This ecological model reproduces the power law of extinctions with $\tau = 2.05 \pm 0.06$, consistent with fossil evidence [5]. These authors also track the role of positive and negative interactions in supporting and destabilizing the ecology, respectively.

There are a number of extensions and close relatives of this model [14–16] and indeed the perspective provided by this and the Bak-Sneppen model has become an essential component of a complete mathematical understanding of the macroevolutionary process [17, 18]. We take from this literature the following stylized facts:

- interactions are an essential component of the evolutionary process;
- these interactions should be generalized away from “food web” pictures;
- interactions should be able to render species more or less viable;
- interacting models tend to exhibit self-organization

An important element that we will not capture in the current model is the role of external driving. In [5] this appears through random changes in γ_{ij} . In [16] the driving is explicitly environmental. For simplicity we exclude these effects. In Section III we will synthesize the stylized facts of Section II into a simple model.

III. THE MODEL

With our model we aim to unify the fitness landscape and ecological approaches to evolutionary dynamics. We recover some qualitative features of the empirical data about macroevolution, including some properties of extinctions and the way genotype space is explored by the population. In construction of the model, we will make frequent reference to the stylized facts that are being implemented at each step.




FIG. 1: Hypercube of dimension 4. Note that each vertex is labeled with a string of 4 0's and 1's. Hence 1111 has neighbors 0111, 1011, 1101, and 1110. The colored sites are elements of percolation clusters.

A. The Hypercube in dimension n

The most primitive choice concerns the structure of the genotype space \mathcal{G} . Since we desire simplicity, we restrict ourselves to a binary genetic alphabet with no further structure (e.g. diploidy). The standard model for \mathcal{G} under these considerations is the hypercube of n dimensions, e.g. [19].

For a natural number n , we construct the graph B^n (n -dimensional hypercube or n -cube) with vertices labeled by all 0, 1 sequences (e_1, e_2, \dots, e_n) of length n ; $e_i \in \{0, 1\}$. Two 0, 1 sequences s_i and s_j are neighbors if they are of Hamming distance $H[s_i, s_j] = 1$, i.e. if they differ by one member of the sequence. For example, 0001100 and 1001100 are neighbors in the graph, and are hence connected by an edge. Note that consequently each vertex has n neighbors and n is the “coordination number” of the n -cube. The total volume (number of sites) of the n -cube is 2^n . In Figure 1 we show a labeled hypercube with coordination number 4.

This geometry models the genotype space of our universe of organisms. Because of the high level of abstraction, we can interpret the sequences variously: as a simple genetic alphabet; as hydrophobic or hydrophilic amino acids in a protein primary structure; or as some kind of binary morphological traits. The neighborhood relations indicate what sequences can be reached via a single point mutation. We do not permit insertions, deletions, or duplications, so the dimension n remains fixed. We also forbid recombination (which would be a type of non-local mutation). Note that even for a relatively short sequence e.g. $n = 50$ the size of the space is enormous (2^{50} possible genomes). Because our fitness function assignment is equivalent to percolation, we will refer to the vertices as “sites” henceforth.




FIG. 2: The Bethe lattice approximation. Here the coordination number is $n = 3$. As $n \rightarrow \infty$ the approximation becomes exact.

B. Fitness function via Percolation

Recall from the stylized facts that selection should matter, but that small differences in fitness should be irrelevant. Further, recall that few genotypes should be fit. Following Gavrillets, we thus independently assign fitness 1 with probability p and fitness 0 with probability $1 - p$ to each site s_i . Denote this assigned value $s_i(1, \cdot)$; respectively $s_i(0, \cdot)$. The notation will become clear as the exposition continues. Assigning a fitness of 0 or 1 is of course a drastic approximation to the actual complexity of the genotype/phenotype map; for example we have introduced no correlation structure on the fitness function (which one might expect *a priori*). The fitness value assigned to each site is called the viability in the literature [11, 12]. The underlying assumption is that an overwhelming majority of sequences are simply non-viable either due to developmental (in the morphological case) or folding errors (in the RNA or protein case), or due to lack of niche (in an ecological perspective). Hence the important distinction is not between the fitness levels of genotypes but rather between those genotypes that are in any way viable (fitness 1) and those that are not (fitness 0).

This fitness function assignment is identical to a percolation process on the n -cube. If $H[s_i(1, \cdot), s_j(1, \cdot)] = 1$ then s_i and s_j are defined to be in the same cluster of fit sequences. A common tactic for dealing with very high dimensional spaces is to approximate the space by a tree with branching number equivalent to the coordination number. This is the so-called Bethe lattice approximation, see Figure 2. The approximation becomes exact as $n \rightarrow \infty$. In the Bethe lattice approximation for ordinary lattices, one takes the limit of infinite lattice size (for lattices of sufficiently high coordination number) and hence expects to see finite *size* corrections for any practical realization. In our case the thermodynamic limit of

infinite size is equivalent to an infinite *dimensional* limit, and one sees finite dimension corrections [21]. It is easy to calculate for the Bethe lattice the critical p at which a “spanning cluster” will appear, $p_c = 1/(n - 1)$. At this p any occupied site will have on average at least one occupied neighbor, so the cluster can persist indefinitely [1, 11, 20]. We also note under the Bethe lattice approximation that [20]:

- the average size of the cluster to which an occupied site belongs scales as $\chi(p) \sim (p_c - p)^{-1}$, $p \rightarrow p_c^-$;
- the characteristic cluster size scales as $s_\xi \sim (p - p_c)^{-2}$, $p \rightarrow p_c^-$;
- the cluster size density scales as $n(s, p) \sim s^{-\tau} \exp(-s/s_\xi)$, $\tau = 5/2$ close to p_c ;
- the “surface area” or number of neighbors of a cluster of size s is $t = 2 + s(n - 2)$.

The scaling of $n(s, p) \sim s^{-5/2}$ is confirmed in our simulations and in studies of percolation on the hypercube in the context of spin glass relaxation [21]. Note that for n quite small the fraction of occupied sites at the percolation threshold is correspondingly small ($1/(n - 1)$), consequently satisfying our stylized fact that fit genotypes should be rare while nevertheless having neutral networks “span” the genotype space.

Here we have adopted the biological interpretation of a cluster as a “neutral network”. See Figure 1 for an example with two percolation clusters. Point mutations from an extant viable genotype allow the occupation of neighboring viable genotypes. We will study the dynamics near the percolation threshold, as this is where the fitness function is interesting. For sub-critical percolation $p < p_c$ we expect many clusters, the largest clusters being of size n [11]. These are relatively loop-free [11, 21] for large n , i.e. they contain few pairs of sites linked by multiple paths within the cluster. Furthermore, for a randomly chosen site in \mathcal{G} , there exists some cluster passing within n steps of that site [11]. For supercritical percolation, $p > p_c$, the largest cluster is of order $2^n/n$, passes arbitrarily close to every site in the sequence space \mathcal{G} , and is typically loopy [11].

In the subcritical case it is easy to interpret a cluster as a “phenotype” i.e. a particular morphology, RNA or protein fold, etc. This is more difficult for the supercritical case without additional genetic complications as studied by Gavrillets [11, 12]. In some empirical neutral networks, the percolating cluster would be naturally partitioned into various phenotype (fold) subclusters, which nevertheless percolate — the subclusters are connected by one mutational step [7, 22, 23].

After setting up the neutral landscape in this fashion, we initialize the model by occupying a randomly chosen, viable site s_0 . For simplicity, we will refer to an occupied site as a “species”, although without a clear notion of biological species (as in [11]) this choice is somewhat arbitrary.

C. Interactions

So far the model construction has followed the fitness landscape tradition. Now we include complications coming from the ecological picture. Namely, we incorporate generalized interactions, capable of rendering sites more or less viable.

With probability q place a dotted directed edge between each ordered pair of distinct sites (s_i, s_j) , $i \neq j$ (note we independently try (s_j, s_i)). For each edge generated, the edge is + with probability 0.5 and – with probability 0.5. Thus we have a directed, signed Erdős-Rényi graph with expected edge number $\langle L \rangle = (2^n)(2^n - 1)q$ and with, on average, half + and half – edges. The mean degree is $\langle k \rangle = \langle L \rangle / (2^n)$, making no distinction on edge direction.

If a site s_i has incident edges (i.e. edges pointing towards it) and the sites from which these edges originate are occupied (call these “activated” edges) we sum the total activated edge symbols with $+ = 1$ and $- = -1$. If the sum is positive, we write $s_i(e, +,)$, where $e \in \{0, 1\}$ is the viability. If the sum is negative we similarly write $s_i(e, -,)$. We interpret these symbols as follows. $s_i(0, +,)$ is *conditionally viable*, i.e. some other species have created a niche for species i (in the molecular RNA/protein interpretation, other molecular species facilitate a fold). $s_i(1, +,)$ is viable. $s_i(1, -,)$ is *conditionally inviable*, meaning that the other species have eliminated the niche via predation, resource destruction, etc. $s_i(0, -,)$ remains inviable. Note that $+, -$ is subject to change as the sources of other incident edges become occupied. Thus a viable site can become viable, then inviable, then viable again depending on the history of the system.

D. Dynamics

We have put extensive structure on our genotype space \mathcal{G} . We can now describe dynamics — an implicit stylized fact of the ecological approach. Start the system by occupying some site $s_0(1, ,)$. Notationally, the occupation of sequence s_0 can be written $s_0(1, , 1)$. This means that s_0 is viable, it has no interactions (because no other sites are occupied) and it is occupied. At every time step, we choose a random occupied site, i.e. extant genotype, choose one of its n neighbors at random and send a mutant to that sequence. Note:

- for realistic (small) mutation rates, the probability of two simultaneous mutations is negligible (thus justifying the choice of one genotype to mutate per time step);
- having a mutant at every time step (as opposed to having some mutation probability) is just a matter of setting the timescale.
- p is very small so most mutations will fail;

If the neighbor s_i is viable or conditionally viable, i.e. $s_i = s_i(1, +, 0)$ or $s_i(0, +, 0)$ then the step is successful (the last element of the triple can be 1 but then no new species is created). We update in this case

- $s_i(1, +, 0) \rightarrow s_i(1, +, 1)$
- $s_i(0, +, 0) \rightarrow s_i(0, +, 1)$.

Otherwise the mutation fails. Note that we allow “back” mutations implicitly. At every occupation, the arrows originating from the newly occupied site become active and influence the sites to which they point. This can cause extinction as well as the creation of conditionally viable sites. We iterate these steps until no further successful mutation is possible. In principle, interactions will create “bridges” across gaps between neutral networks, allowing major “speciation” events to occur.

We illustrate the dynamics in Figure 3. Mutation explores the upper cluster (red) until a conditionally viable site (gray with white boundary) is created linking the upper and lower clusters. This site is then occupied and the lower cluster explored, rendering an element of the upper cluster conditionally inviable (white with gray boundary).

E. Numerics

Because of the small value of p — that is, the stylized fact that few genotypes should be viable — most mutations will be failures. If we were to simulate the dynamics directly, as described above, most computer time would be wasted on unsuccessful jumps to inviable sites. Consequently, we use an event-driven Gillespie scheme to select the next successful mutation and randomly generate an appropriate waiting time [24].

Specifically, every potential mutation from an occupied site s_i to a viable but unoccupied site s_j is assigned an equal propensity $pr(i, j) = 1$. All such ordered pairs have an equal probability of being selected for implementation. After a pair is selected, we generate the amount of simulation time $t(i, j)$ that passes before the event occurs. Let X be a random number chosen from the uniform distribution over $[0, 1)$, and $T = \sum pr(i, j)$, i.e. the number of potential viable mutation steps. Then the waiting time $t(i, j) = \frac{1}{T} \log\left(\frac{1}{X}\right)$ [24]. This approach conserves computing time by assuring that every simulation step involves a successful mutation.

IV. RESULTS

Our model is parameterized by the dimension of the genotype space, n ; the percolation probability, p ; and the interaction probability q . It is obviously desirable to study \mathcal{G} with as high dimension as possible. With available computational resources we were able to simulate $n = 50$. Each $n = 50$ simulation lasted approximately

70,000 simulation steps, the limiting factor being memory usage.

We selected values of p close to the critical value $p_c = 1/(n - 1) = 1/49 \approx 0.02041$. The choice of just sub- or super-critical p appeared to have little impact on model outcome for fixed values of q . The selection of q was slightly more difficult. For subcritical p we scanned the parameter space of q near $q = 1/n^2$, evaluating the probability that a simulation runs for its full lifetime (strictly speaking, we used as an indicator the successful occupation of n^2 sites — much larger than any initial cluster size — suggesting that the simulation has escaped the initial cluster with overwhelming probability). This “escape probability” as a function of q is plotted in Figure 4. We find that indeed near $q = 1/n^2 = .0004$ the escape probability is non-zero. Most simulations were conducted for $q = 2/n^2$, where we expect between 10% and 20% of simulations to “escape”.

How are sites occupied in the model over the course of a simulation run? It is conceivable, for example, that most activity occurs “near” the originating site, with sites continually discovered, rendered inviable, and then reoccupied when they become viable again. But keeping track of the number of realized sites, (i.e. sites that are *ever* occupied in the course of a simulation) and the number of extant sites (i.e. sites occupied at a particular simulation step) reveals a very different picture. In Figure 5 we see that the total number of realized sites (upper green line) increases almost linearly with the number of simulation steps. Hence at almost every simulation step the unfolding biosphere is discovering new sites or genotypes, rather than revisiting occupied territory. The fact that the number of extant species (lower blue line) fluctuates around 2000 species over the course of the simulation indicates that there is an advancing front of activity, which explores genotype space while eliminating older genotypes as it advances. This behavior is typical of simulations that escape the initial cluster.

Rather surprisingly, this picture of an advancing front is consistent with the morphological diversification of blastozoans discussed by Gavrillets [25]. In this analysis the species are scored on n characters. Gavrillets then measures the average morphological disparity within extant species, the taxonomic diversity, and the average morphological distance from the founder species. The picture described by Gavrillets of a “compact group” moving away from the founder species in morphospace is extremely similar to the interpretation offered here of our simulation results. We can test the similarity of the two pictures by explicitly calculating at each simulation step the average Hamming distance from the founding site (among extant, occupied sites) and the average pairwise Hamming distance between extant species. Gavrillets finds for the fossil data that the former increases while the latter grows initially but then shrinks and stabilizes at a small value. We will measure these quantities in future work, and predict behavior very similar to the fossil data.




FIG. 3: Model Dynamics for $n = 4$. Note the interactions rendering sites viable (thus connecting the upper and lower clusters) and inviable (eliminating a member of the upper cluster).

We now turn to the role of interactions in facilitating the initial exploration of genotype space and subsequent stability at roughly 2000 extant species. Consider all pairs of viable sites (i, j) . Let the number of such pairs at time T be $N(T)$. Define a time dependent quantity $f_T(i, j)$ taking the value -1 if the interaction is negative at time T and $+1$ if the interaction is positive (recall that when set up, roughly half of the edges are positive and half negative; but we restrict here to edges between viable sites). Then the interaction bias, summing over all pairs (i, j) :

$$I_{av}(T) = \frac{1}{N(T)} \sum_{i,j} f_T(i, j) \quad (2)$$

provides a measure of the overall bias of interaction towards viability. If $I_{av}(T) = 0$ then positive and negative interactions are perfectly balanced, and indeed absent any dynamics this is what we should expect, on average.

In Figure 6 we see a typical time series (in fact the same time series from Figure 5) against which we also plot the interaction bias as a function of simulation time T (rather than simulation steps). We have focused on the “interesting region” after the simulation has grown towards its typical number of occupied sites. The period of growth is associated with a strong positive interaction bias, while the stabilization of the simulation into a

steady state fluctuating around 2000 occupied sites correlates with a decrease in interaction bias towards a small but nevertheless positive bias. This indicates that the steady state is “supported” by a bias towards positive interactions, while there are simultaneously enough interactions and enough occupied sites that a substantial number of negative interactions can take place between *viable* sites without rendering those sites extinct. There seems to be some correlation between increases in the interaction bias and increases in the number of extant species, and likewise between decreases in the bias and periods of sequential extinction, though this remains to be quantified. The apparent correlation is somewhat similar to the relationship between interaction entropy (similar to our interaction bias) and extinction measured in [5].

In order to understand the behavior of the time series, we performed a detrended fluctuation analysis [26]. We briefly summarize this technique. The data are divided into windows of length L . Within each window, we calculate a best-fit trend, in this case a polynomial of degree 2. Then all fluctuations around the trend are summed for each window, normalized by the length of the window, and summed in turn, giving the total fluctuation $F(L)$. Plotting $F(L)$ against L gives a power-law fit for self-similar time series, $F(L) = L^\alpha$. For our data we estimate $\alpha = 1.5065$ for simulations on a hypercube




FIG. 4: In this plot we show the escape probability as a function of the interaction parameter q , which sets the likelihood for any ordered pair of sites to have an interaction. $n = 50$ and $p = 0.02037$, i.e. p is slightly sub-critical. The first non-zero point on the plot is very close to $q = 2/n^2$, at which value most simulations were performed.




FIG. 6: The relationship between number of species and interaction bias. The blue line (upper after $T = 8.0$, left axis) shows the number of occupied sites (extant species) as a function of simulation time T . The red line (lower, right axis) shows the interaction bias. Note that the interaction bias remains relatively constant as the simulation settles into the apparent steady state. Model parameters are $n = 50$, $p = .02037$ (slightly subcritical) and $q = .0008 = 2/n^2$.

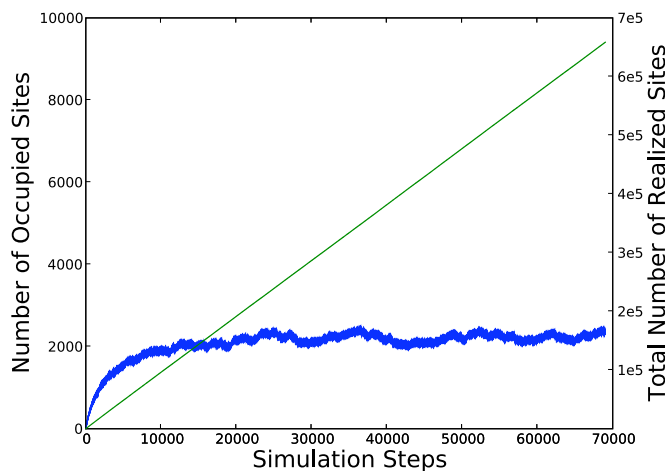


FIG. 5: The lower, blue line (left axis) shows the number of occupied sites (extant “species”) as a function of simulation steps. Note that after an initial period of growth this fluctuates around 2000 occupied sites at any given simulation step. The upper, green line (right axis) shows the number of realized sites, i.e. sites ever occupied in the course of the simulation, as a function of simulation steps. This quantity increases almost linearly with the number of simulation steps, indicating that evolution is continually exploring new parts of genotype space. Model parameters are $n = 50$, $p = .02037$ (slightly subcritical) and $q = .0008 = 2/n^2$.

of dimension $n = 50$, see Figure 7. The fit of $\alpha \approx 3/2$ works well for $n = 40$ and $n = 30$ as shown in the Figure. $\alpha = 3/2$ indicates that the correlations in the time series are similar to those observed in Brownian motion. Generally, the time series exhibits fractal scaling insensitive to the dimension of the underlying genotype space. Comparison with detrended fluctuation analysis of taxonomic diversity data such as that from [25] will be a fruitful direction for future work and a direct test of the model.

We now consider the time series of the extinction rate, as observed in simulations that survive to full simulation time. Dividing the simulation time into suitably coarse-grained periods of length 1000 (to emulate the division of fossil data into stratigraphical periods), we consider the ratio of the number of species that went extinct during that time period to the number of species that existed during that time period. This provides an estimate of the probability that a randomly selected, living species (or occupied site) will go extinct during the time period. We show a typical plot of this time series for subcritical p (using the same $n = 50$ data as in previous plots) in Figure 8. Comparison with typical fossil data, e.g. Figure 2(A) of [4] shows strong *qualitative* similarity but a much smaller variation in amplitude. Our simulations typically show somewhere between 25% and 42% extinction, while the fossil data varies from close to 0% up to 50%. It seems that the interactions of the model are unable to generate really catastrophic biodiversity loss, while paradoxically also maintaining a high background rate of extinction.




FIG. 7: The detrended fluctuation analysis for three typical simulation runs. Red dots are for a genospace of dimension $n = 30$, blue dots are $n = 40$, and green dots are $n = 50$. p is slightly subcritical, $p = 1/(n - .9)$ and $q = 2/n^2$ for the appropriate values of n . Note the logarithmic axes; all three plots are well fit by a power law with $\alpha \approx 3/2$.




FIG. 8: Extinction rate averaged over periods of length 1000. Parameters are $n = 50$, $p = .02037$ (slightly subcritical) and $q = .0008 = 2/n^2$. Note the high background rate of extinction, i.e. no real periods of stasis.

In other words, we do not see periods of “stasis”, i.e. comparably low extinction rate, as in (for example) the Bak-Sneppen model [4].

Recall that when a new site is occupied (i.e. a new “species” appears) the interactions thereby created may make other sites inviable, i.e. cause other species to go extinct. If we define the number of sites rendered inviable as S , and consider each such instance an “avalanche”,




FIG. 9: $N(S)$, the number of extinctions of size S , for three typical simulations. Extinctions are caused by the occupation of a new site at each time step, and their size is defined as the number of species that go extinct in that time step. Red dots are for a genospace of dimension $n = 30$, blue dots are $n = 40$, and green dots are $n = 50$. p is slightly subcritical, $p = 1/(n - .9)$ and $q = 2/n^2$ for the appropriate values of n . The distribution is fat-tailed but spans too few orders of magnitude for a reliable power-law fit. The distribution is insensitive to changes in n and small changes in p and q , including supercritical p (not shown).

we can measure the size distribution $N(S)$ of extinction events. Several typical distributions are plotted in Figure 9. Note that the distributions are essentially identical despite variation in the size of the underlying hypercube ($n = 50, n = 40, n = 30$, simulations used in Figure 7). Although our data span only one and a half orders of magnitude, they seem to be fat-tailed and one could fit a power law $N(S) \sim S^{-\alpha}$ with α between 2.5 and 3.0. Such power law scaling is typical of ecological models [4, 5, 13, 18] and while we come as close to the empirical value $\alpha = 2.0$ as Bak-Sneppen type models, other models [5, 16] generate exponents almost identical to the empirical value. Our result is unchanged by any suitable generalization of avalanche causality, and appears insensitive to changes in n as well as small changes in p (including supercritical) and q . It is perhaps a result of the memory limitations of our model, which forbid the very long runs where one might expect to see large extinctions typical of most self-organized critical models. But the absence of a clean power-law or very large events indicates that our model may not be a case of self-organized criticality.

Further evidence that our model is not an example of self-organized criticality comes from consideration of the waiting time T between extinction events. T is defined as the simulation time that passes between two extinctions events. For self-organized critical models this distribution is typically power law, i.e. $N(T) \sim T^{-\gamma}$ for

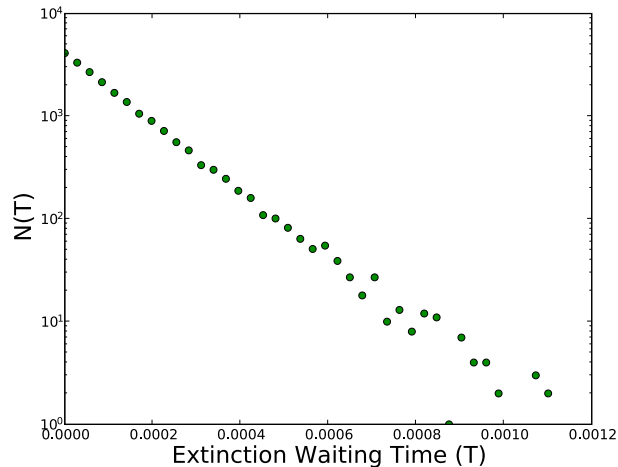


FIG. 10: The distribution of waiting times, T between extinction events. The distribution is almost certainly exponential (straight line on a log-linear plot). Parameters are $n = 50$, $p = .02037$ (slightly subcritical) and $q = .0008 = 2/n^2$. The exponential distribution is insensitive to changes in p and q . Were our model an instance of self-organized criticality, the distribution would be power law [14].

$\gamma = 3.0 \pm 0.1$ as in [14]. A typical distribution for a simulation of our model, shown in Figure 10, is by contrast exponential. This feature appears robust against small variations in p and q (and measuring T in simulation steps). We have some evidence of n insensitivity as similar results hold for $n = 20$. We should also note that the interpretation of extinctions as a self-organized critical phenomenon has been questioned [27].

V. DISCUSSION, INTERPRETATION, AND FUTURE WORK

A. Discussion

Our model was designed to encapsulate the basic stylized facts about evolution encoded in the fitness landscape and ecological approaches. We believe each of these approaches provides an important perspective on the nature of selection and the drivers of macroevolution. In a sense, our model remedies important defects in each approach. To ecological models it adds the notion of an underlying genospace, whose topology, induced by the fitness landscape, guides the speciation driven by the ecological dynamics. To fitness landscape models it adds explicit notions of coevolutionary dynamics and the construction of niches by other species. In fact, our model illustrates that one need not impose percolation of a neutral network by hand, setting a supercritical p as in [1, 11, 12]. Rather, ecological interactions create a context in which bridges can be built between the neu-

tral networks of subcritical percolation. As a first step towards constructing a coherent model of evolutionary dynamics — one that honors most underlying qualitative features without sacrificing simplicity — our model is an important contribution and starting place.

Turning to the results, we note that our exploration of parameter space has been somewhat limited, concentrating on near critical p and values of q for which escape probability first becomes non-zero. Nevertheless, we find a satisfying robustness of many statistical features to variation of these parameters, as well as the dimension of \mathcal{G} . Furthermore, we rather unexpectedly observed a movement of extant species through genospace (the “advancing front”) in striking similarity to the empirical results for blastozoans [25]. This feature is even more surprising when we consider that the model generating similar behavior in [25] not only omits the fitness landscape but also uses exogenous events rather than endogenous dynamics to drive extinctions. Further exploration of this similarity is a major goal of further research.

Our model also self-organizes a system of positive, self-supporting interactions, stabilizing the number of extant species. In other words, we observed the emergence of a sort of ecology, where the various species create and sustain niches for one another. However, we need to explore the correlations between overall interaction bias and extinction to understand this niche construction mechanism more fully.

The time series of species observed in our simulations displays self-similarity consistent with Brownian motion. At the moment it is rather unclear how to interpret this result. An obvious first step would be to examine e.g. the blastozoan data [25] and subject it to detrended fluctuation analysis. The robustness of the Brownian motion result indicates that this is probably an inevitable consequence of the current model, so supplying paleontological or empirical evidence of such behavior elsewhere is quite important.

We finally note that our model produced some but not all features of self-organized criticality models of evolution [4, 5, 16, 18]. While we generate time series of extinction rates that are superficially similar to empirical time series, they display limited range and very high background extinction rates. We find fat-tailed and possibly scale free distributions of extinction size, but do not find the very large events characteristic of the best ecological models or the fossil data. This perhaps follows from current limitations on simulation run time, which may be overcome in the future. Most crucially, we do not find power law distributed waiting times, which suggest that our model is not an elaborate form of self-organized criticality. This result in itself is quite interesting, as one might wonder how many empirical features of macroevolution can be reproduced *without* a self-organized critical model (although again see [27]).

In the next subsection we turn to possible interpretations of the model. While we have followed a concrete, ecological interpretation throughout the exposition, we

believe the rich set of possible interpretations illustrates the power of our approach.

B. Interpretation

The model proposed in this paper is highly abstract. While this abstraction complicates direct comparison with data, it facilitates a wide range of interpretation and hence application. The interpretation guiding the development and exposition of the model was explicitly macroevolutionary and ecological. The underlying sequence space \mathcal{G} in this case is interpreted morphologically [25], and each site can reasonably be interpreted as a new “species” (in which case clusters would be interpreted as some higher taxonomical unit). The fitness function measures the evolutionary viability of a particular collection of morphological traits, i.e. the traits must be achievable developmentally and they must fit into an ecological niche. The interactions further shape these ecological niches. As an example, consider an expensive morphological innovation like a heavy beak, which a predatory bird might use for cracking open the shells of turtles. In the absence of prey requiring the innovation, heavy-beaked mutants would be at an extreme disadvantage and disappear: the cost of the beak is not worth its benefit. But if a shelled creature were to emerge, invulnerable to extant predators — but vulnerable to the shell-cracking “shredder” beak of this mutant — then strong selective pressure would favor the appearance of such large-billed predators. The “heavy-beak” mutation might in turn open up a new area of sequence space (i.e. a new cluster) to these birds, providing developmental raw material which could be shaped by mutation into a variety of specialized morphologies.

To illustrate the flexibility of this model, we provide three alternate interpretations, one in terms of protein evolution, one in terms of RNA evolution, and one in terms of technological innovation. This in no way exhausts the space of interpretations; one can also frame the model as a sort of autocatalytic set of chemical reactions, as a self-assembly process, etc. One can even formulate a spin-glass-like model closely related to our model, in which relaxation is localized.

1. Protein Evolution

If we interpret the model as a description of protein evolution, each site in \mathcal{G} represents a protein sequence of a given length n (the coordination number of the hypercube). The genotype space (better, sequence space) contains all possible protein sequences of that length. In full generality this would be 20^n sites, as there are 20 amino acids. As we set the linear dimension of the hypercube at 2 (a binary hypercube) we implicitly assume two possible amino acids (or a coarse-graining into hydrophobic and hydrophilic amino acids). A cluster of a

priori viable sites is a web of neutral neighbors — sequences that differ, but that share the same basic three-dimensional structure or fold. We call these neutral because it is generally the three-dimensional structure of a protein, as opposed to its specific sequence, that determines protein function. This illustrates the degeneracy of the genotype-phenotype map. Three-dimensional protein structures are generally far more conserved through evolution than amino acid sequences [28]. Many sequences with no detectable sequence similarity share essentially the same structure, and presumably, function [29].

In the current version of the model, we only allow forces internal to the model to change the viability of sites. The “universe” only includes proteins of a given chain length, so changes to viability occur only through proteins of this single length. We do not consider changes to the viability of sites that might occur through forces not represented in the genotype space — proteins of different lengths, environmental changes, etc.

On this interpretation of the model, viable sites are those that can fold without the help of chaperone proteins (at least those in the model “universe”, i.e. those of length n) into a stable minimum free energy state. Recall that conditionally viable sites can be occupied if sites with + arrows pointing into the conditionally viable site are also occupied. The “activation” of a site with a + arrow is like the evolution of a chaperone protein (of length n) that allows a potentially unfoldable protein (of length n) to fold and hence become viable. Chaperones prevent protein aggregation by binding polypeptide folding intermediates as soon as they emerge from the ribosomal exit tunnel, thus playing a crucial role in the creation of functional proteins with well-defined, three-dimensional conformations [30]. Such intermolecular interactions are strongly favored in a crowded cellular environment where numerous unfolded polypeptides are translated in close proximity to one another.

Sites that are conditionally inviable due to – arrows coming into them can be thought of as proteins adversely affected by mutant copies acting in a dominant negative manner. This kind of mutant protein interferes with the activity of an otherwise functional, normal copy of the protein— for example, via competitive inhibition (where a non-functional protein can still bind a target, blocking the normal protein from binding) or dimerization (where non-functional copies of the protein combine with functional copies).

An obvious extension of the model would allow node viability to be changed by forces that are external to the explicitly modeled genotype space. Anything affecting protein viability that is not another protein of the same length n would count as external — e.g. developmental and environmental changes that open up new realms of possibility for protein evolution.

It would also be interesting to allow two sites to recombine, thus permitting a jump across sequence space. This extension is especially important in the context of protein evolution because innovation via recombination

may well be the primary means by which proteins explore new structures (i.e. clusters of viable sites) [31–34].

2. RNA Evolution

Interpreting our model in terms of RNA evolution is even more straightforward than the protein case. For in the protein case we suppress the actual genetic code underwriting the amino acid sequences, dealing only with the space of all amino acid sequences (although it is the genetic code that undergoes mutation). In the RNA case, the RNA sequence itself both undergoes mutation and folds into the structures that are acted upon by evolution. Another way of putting this: in the RNA case, the genotype-phenotype map is quite direct [7].

In reality RNA sequences are specified over a four element alphabet $\{A, U, G, C\}$ so that the genotype space for sequences of length n would contain 4^n sites. There are several possible levels of resolution at which RNA structure can be specified; the best compromise between theoretical tractability and empirical accessibility is at the level of secondary structure [22]. So in our model viable sites represent RNA sequences that fold reliably into some “fit” secondary structure and clusters are neutral networks of sequences all folding into the same secondary structure. It is known from empirical, numerical, and mathematical work that the neutral networks corresponding to any pair of different secondary structures almost touch, and indeed can be found within a few point mutations of an arbitrarily chosen point of \mathcal{G} — the so called shape space covering conjecture [7, 35]. Thus the picture from our model of neutral networks nearly percolating sequence space is quite realistic in the RNA case.

RNA-RNA interaction is by now a well-known phenomenon, underwriting RNA interference, microRNA-messengerRNA binding, antisense interactions, etc. Interaction is a crucial feature of our model, and – interactions could be interpreted as instances of RNA interference or competitive binding, while + interactions could represent RNA chaperones, of which there are empirical examples [36]. RNA interaction can be modeled directly with contemporary thermodynamic and folding algorithms [37, 38]. This means that in fact we could quite easily study realistic coevolution of short RNA (e.g. microRNA of 19-25 nucleotides) by supplementing our model with these techniques. As in the protein case, our limiting of the mutation operator to point mutations, (excluding insertions, deletions, duplications, and recombinations) is a substantial limitation and simplification.

3. Technological Innovation

Departing from the biological context of prior interpretations, we now offer an economic interpretation of the model. Since the early work of Schumpeter [39] and von Hayek [40], the tradition of evolutionary economics

[41, 42] has emphasized close similarities between evolutionary biology and economic development. We follow this intuition in attempting this third interpretation.

Consider a product defined by n binary characteristics. The space of all 0, 1 sequences of length n then defines \mathcal{G} , the space of all possible products. We assume that innovations may occur only by altering one particular product trait (e.g. developing a teddy bear with posable limbs from one with “floppy” limbs) and that the binary economic viability (something like the *a priori* demand for the good) is independently, identically distributed. Viable sites under the percolation process are “in demand”, in the absence of other goods, and would be successful if developed.

Interactions are similarly interpreted as altering the market niche of the good in question. For example, the development of posable teddy bears might create a niche for teddy bear furniture, which would be useless in the context of “floppy” bears. Products can also weaken or destroy the viability of other goods, as for example next generation game systems tend to do for their predecessor systems.

At the moment this interpretation merely illustrates the flexibility of our model (or alternately its extreme abstraction). It is not too difficult to imagine, however, experimental economics investigating the evolution of n binary trait goods. One could also try to collect economic data similar to the blastozoan data used in morphological evolution models and validate or reject the general set-up of this economic model with such data [25].

C. Future Work

While investigation of ecological models of macroevolution has slowed considerably, interest in holey fitness landscapes remains high. Indeed the framework provided by our extension of the holey fitness landscape idea may be instrumental in understanding e.g. recent work on epistasis and evolvability [43]. We thus outline several promising avenues for future work. These divide broadly into work providing better understanding of the present model and work extending the model in more realistic directions.

Most immediately, a more thorough characterization of the parameter space of the present model is necessary. This is an area where analytic work would be of great help, as it would be free of the memory limitations hindering numerical work. Thus far our exploration of the model has been exclusively numerical, despite the success of analytic work on holey fitness landscapes [1, 11, 12, 35]. Even more interesting is the possibility of empirically testing predictions of the model, for example the comparison of the time-dependent geometry of the set of extant species with fossil data [25]. Another fruitful line of investigation would involve the definition of plausible phylogenetic trees based on simulations of the model. Many statistical features of empirical phylogenetic trees

cannot be reproduced by simple models and perhaps the structuring of genotype space by fitness landscapes and coevolutionary interaction is precisely the missing ingredient [44]. And comparison with in vitro RNA coevolution models would provide an extraordinary opportunity to put the model to direct test. From a more mathematical perspective, modern computational homology techniques could provide considerable insight into the geometry of holey fitness landscapes and their descendants in our model, rare examples of complicated high dimensional spaces with immediate empirical relevance [1].

Turning to realistic extensions of the model, an obvious first step would be adding external driving in the form of exogenous extinction. One could also imagine extending the set of mutation operators so that n could change over the course of a simulation run. Thus we could incorporate insertions, deletions, and duplications, as well as recombinations. This would enhance the realism of mutation in the model considerably, at the cost of numerical complication. Another realistic modification would copy the interaction structure of a mutant largely from the parent, with some small variation. Indeed inheritance of interaction is crucial in ecological models for

the self-organization of large extinction events [5]. Even more realistic, but computationally more difficult, would be context dependencies for the activation of outgoing interactions from a viable site.

The ease with which we can identify directions for further research and plausible extensions of the model illustrates the still largely untapped potential of unifying the two major metaphors in evolutionary modeling: fitness landscapes and ecology. We believe our model provides an important first step into this as yet unexplored terrain, and look forward to much fruitful work following in these tentative footsteps.

Acknowledgements: This work was partially supported by the Santa Fe Institute whose research and education programs are supported by core funding from the National Science Foundation and by gifts and grants from individuals, corporations, other foundations, and members of the Institute's Business Network for Complex Systems Research. The authors gratefully acknowledge stimulating conversations with Sarah Cobey, Devin Drown, Walter Fontana, and D. Eric Smith. JGF thanks iCORE for its generous support.

-
- [1] S. Gavrillets, Evolution and Speciation in a Hyperspace: The Roles of Neutrality, Selection, Mutation, and Random Drift. In J.P. Crutchfield and P. Schuster, eds., *Evolutionary Dynamics: Exploring the Interplay of Selection, Accident, Neutrality, and Function* (Oxford University Press, Oxford, 2003).
- [2] S. Wright, The roles of mutation, inbreeding, crossbreeding and selection in evolution. In D.F. Jones, ed., *Proceedings of the Sixth International Congress on Genetics* **1**, 356-366 (1932).
- [3] J. Hofbauer and K. Sigmund, *Evolutionary Games and Population Dynamics* (Cambridge University Press, 1998).
- [4] K. Sneppen, P. Bak, H. Flyvbjerg, and M.H. Jensen, Evolution as a self-organized critical phenomenon. *Proc. Natl. Acad. Sci. USA* **92**, 5209-5213 (1995).
- [5] R.V. Solé, J. Bascompte, and S.C. Manrubia, Extinction: Bad Genes or Weak Chaos? *Proc. R. Soc. Lond. B* **263**, No. 1375 (1996).
- [6] S.J. Gould, *Wonderful Life* (W.W. Norton and Co., 1989)
- [7] B. M. R. Stadler, P.F. Stadler, G.P. Wagner, and W. Fontana, The Topology of the Possible: Formal Spaces Underlying Patterns of Evolutionary Change, *J. theor. Biol.* **213**, 241-274 (2001).
- [8] S. Gavrillets, Evolutionary Ecology: Fitness Landscapes. In S.E. Jorgensen and B. Fath, eds., *Encyclopedia of Ecology* (Elsevier, 2008).
- [9] M. Kimura, *The neutral theory of molecular evolution* (Cambridge University Press, New York, 1983).
- [10] J. Maynard Smith, Natural Selection and the Concept of a Protein Space. *Nature* **255**, 563-564 (1970).
- [11] S. Gavrillets and J. Gravner, Percolation on the Fitness Hypercube and the Evolution of Reproductive Isolation. *J. theor. Biol.* **184**, 51-64 (1997).
- [12] J. Gravner, D. Pitman, S. Gavrillets, Percolation on fitness landscapes: Effects of correlation, phenotype, and incompatibilities. *J. theor. Biol.* **248**, 627-645 (2007).
- [13] P. Bak and K. Sneppen, Punctuated equilibrium and criticality in a simple model of evolution. *Phys. Rev. Lett.* **71**, 4087-4090 (1993).
- [14] R.V. Solé and S.C. Manrubia, Extinction and self-organized criticality in a model of large-scale evolution. *Phys. Rev E* **42**, 1 (1996).
- [15] S.C. Manrubia, M. Paczuski, A Simple Model of Large Scale Organization in Evolution. *Int. J. of Mod. Phys. C* **9**, 1025-1032 (1998).
- [16] B.W. Roberts and M.E.J. Newman, A model for evolution and extinction. *J. theor. Biol.* **180**, 39-54 (1996).
- [17] P. Bak and M. Paczuski, Complexity, contingency, and criticality. *Proc. Natl. Acad. Sci. USA* **92**, 6689-6696 (1995).
- [18] R.V. Solé, S.C. Manrubia, M. Benton, S. Kauffman, and P. Bak, Criticality and scaling in evolutionary ecology. *Trends in Evolution and Ecology* **14**, 4 (1999).
- [19] M. Eigen, R. Winkler-Oswatitsch, A. Dress, Statistical geometry in sequence space: a method of quantitative comparative sequence analysis. *Proc. Natl. Acad. Sci. USA* **85**, 5913-5917 (1988).
- [20] K. Christensen, N.R. Moloney, *Complexity and Criticality* (Imperial College Press, London, 2005).
- [21] B. Kahng, Percolation in the hypercube and the Ising spin-glass relaxation. *Phys. Rev. A* **43**, 1791-1801 (1991).
- [22] W. Fontana, Modelling 'evo-devo' with RNA. *Bioessays* **24**, 1164-1177 (2002).
- [23] A. Babajide, R. Farber, I.L. Hofacker, J. Inman, A.S. Lapedes, and P.F. Stadler, Exploring Protein Sequence Space Using Knowledge-based Potentials, *J. theor. Biol.* **212** 35-46 (2001).

- [24] D.T. Gillespie, Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.* **81**(25), 2340-2361 (1977).
- [25] S. Gavrillets, Dynamics of clade diversification on the morphological hypercube. *Proc. R. Soc. Lond. B* **266**, 817-824 (1999).
- [26] C.K. Peng, S.V. Buldyrev, S. Havlin, M. Simons, H.E. Stanley, A.L. Goldberger, Mosaic organization of DNA nucleotides. *Phys. Rev. E* **49**, 1685-1689 (1994).
- [27] J.W. Kirchner and A. Weil, Correlations in fossil extinction and origination rates through geological time. *Proc. Roy. Soc. Lond. B* **267**, 1301-1309 (2000).
- [28] C. Chothia and A.M. Lesk, The relation between the divergence of sequence and structure in proteins. *EMBO J.* **5**, 823-826 (1986).
- [29] L. Lo Cone, B. Ailey, T.J. Hubbard, S.E. Brenner, A.G. Murzin and C. Chothia, SCOP: A structural classification of proteins database. *Nucleic Acids Research* **28**, 257-259 (2000).
- [30] J.C. Young, V.R. Agashe, K. Siegers, and F.U. Hartl, Pathways of chaperone-mediated protein folding in cytosol. *Nature Reviews Molecular Cell Biology* **5**, 781-791 (2004).
- [31] L. Patthy, Introns and exons. *Current Opinion in Structural Biology* **4**, 383-392 (1994).
- [32] Y. Xia and M. Levitt, Roles of mutation and recombination in the evolution of protein thermodynamics. *Proc. Natl. Acad. Sci. USA* **99**, 10382-10387 (2002).
- [33] L.D. Bogarad and M.W. Deem, A hierarchical approach to protein molecular evolution. *Proc. Natl. Acad. Sci. USA* **96**, 2591-2595 (1999).
- [34] Y. Cui, W.H. Wong, E. Bornberg-Bauer and H.S. Chan, Recombinatoric exploration of novel folded structures: A heteropolymer-based model of protein evolutionary landscapes. *Proc. Natl. Acad. Sci. USA* **99**(2), 809-814 (2002).
- [35] C. Reidys, P. Stadler, P. Schuster, Generic properties of combinatory maps: neutral networks of RNA secondary structures. *Bull. Math. Biol.* **59**, 339-397 (1997).
- [36] T. Geissmann and D. Touati, Hfq, a new chaperone role: binding to messenger RNA determines access for small RNA regulator. *EMBO J.* **23**, 396-405 (2004).
- [37] U. Muckstein, H. Tafer, J. Hackermüller, S. Bernhart, P. Stadler and I. Hofacker, Thermodynamics of RNA-RNA binding. *Bioinformatics* **22**, 1177-1182 (2006).
- [38] S. Bernhart, H. Tafer, U. Muckstein, C. Flamm, P. Stader and I. Hofacker, Partition function and base pairing probabilities of RNA heterodimers. *Algorithms Mol. Biol.* **1**, 3 (2006).
- [39] J. Schumpeter, *The Theory of Economic Development* (Harvard University Press, 1934).
- [40] F.A. Hayek, *The Constitution of Liberty* (University of Chicago Press, 1978).
- [41] R.R. Nelson and S.G. Winter, *An Evolutionary Theory of Economic Change* (Harvard University Press, 1982).
- [42] E.D. Beinhocker, *The Origin of Wealth: Evolution, Complexity, and the Radical Reworking of Economics* (Harvard Business School Press, 2006).
- [43] A. Le Rouzic, O. Carlborg, *Evolutionary potential of hidden genetic variation. Trends in Ecology and Evolution* **23**, 33-37 (2008).
- [44] E.A. Herrada, C.J. Tessone, K. Klemm, V.M. Equiluz, E. Hernández-García, and C.M. Duarte, Universal Scaling in the Branching of the Tree of Life. *PLoS One*, **3**, e2757 (2008).