

Chapter 26

The Challenge of Scale in Molecular Adaptation: Local Searches in Astronomical Genotype Networks



Susanna Manrubia, Luis F. Seoane, and José A. Cuesta

Abstract The exploration of vast genotype spaces poses fundamental challenges for evolving populations. As the number of genotypes encoding viable phenotypes grows exponentially with genome length, populations can only explore a tiny fraction of these immense spaces, a fact consistently supported by empirical and theoretical evidence. Paradoxically, local, mutation-driven searches near abundant sequences allow populations to generate phenotypic improvements and functional innovations despite this immense search space. In this contribution, we integrate insights from viral evolution with theoretical expectations derived from genotype–phenotype maps to reexamine how high-dimensional sequence spaces shape evolutionary dynamics. In resolving the paradox, abundant phenotypes play a crucial role because their combinatorial weight biases evolutionary trajectories. We discuss how this bias, together with limited accessibility of fitness peaks, modifies traditional metaphors—such as fitness landscapes—and challenges standard notions of evolutionary optimality. Our results underscore that adaptation is predominantly local yet remarkably efficient, providing a unifying perspective on the coexistence of robustness, innovation, and constrained exploration in molecular evolution.

S. Manrubia (✉)

Museo Nacional de Ciencias Naturales (CSIC), Madrid, Spain

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

e-mail: smanrubia@mncn.csic.es

L. F. Seoane

Institut de Biologia Evolutiva (UPF-CSIC), Barcelona, Spain

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

e-mail: lf.seoane@ibe.upf-csic.es

J. A. Cuesta

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

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

e-mail: cuesta@math.uc3m.es

1 Introduction

Genetic diversity is a fundamental prerequisite for species survival and evolution (Agashe et al., 2023). Natural populations produce and maintain high levels of diversity through multiple mechanisms, ensuring efficient exploration of functional alternatives and rapid responses to changing environments. Our knowledge of intraspecies diversity is particularly extensive for viruses and microbial populations. High-throughput sequencing and improved sampling have repeatedly revealed substantial genetic diversity within single hosts (Lauring, 2020), within microbial species in complex environments, and across replicate experimental populations. Deep sequencing of SARS-CoV-2 (Jones, 2025) and other RNA viruses (Bordería et al., 2011) has identified measurable intra-host single-nucleotide variants (Seoane et al., 2026), persistent mixed infections in some patients, and lineage-specific differences in within-host diversity (Gu et al., 2023). These studies and many others indicate that viral populations commonly exist as diverse “clouds” of related genotypes consistent with a quasispecies-like structure (Domingo, 2006).

Despite the typically lower mutation rates of prokaryotes and smaller effective population sizes compared with viral populations, metagenomics and genome-resolved sequencing have revealed high standing diversity within bacterial species in natural microbiomes (Kim et al., 2024). Many species are not represented by a single dominant genotype but by coexisting strain lineages and within-strain polymorphisms (Viver et al., 2024). Long-read and single-cell approaches uncover fine-scale structures—haplotypes, mobile elements, and recombination tracts—that substantially increase inferred within-taxon diversity relative to earlier short-read-based snapshots. These data indicate that bacteria *in situ* maintain far more genetic variation than simple single-clone models predicted (Harris et al., 2021).

In vivo observations are complemented by *in vitro* evolution experiments, which show that, in large asexual populations, adaptation often involves multiple coexisting beneficial lineages (Lang et al., 2013), polygenic shifts, and frequent soft sweeps (Harris et al., 2021). Populations experience both rapid selective sweeps and extended periods of balanced polymorphism or competing mutations (Laguna-Castro et al., 2025), depending on population size, mutation supply, recombination, and ecological structure. These findings suggest that periods away from mutation–selection equilibrium may be the natural state of microbial (Cvijović et al., 2018) and viral (Braun et al., 2018) populations.

Several theoretical frameworks have been developed to explain these observations. Kimura’s Neutral Theory (Kimura, 1968) proposed that most molecular substitutions are effectively neutral, with substitution rates equal to mutation rates—a powerful null model for sequence divergence. Ohta extended this to the nearly neutral regime (Ohta, 1973), emphasizing that the fate of weakly selected mutations depends on population size and drift. Maynard Smith’s protein space metaphor (Maynard Smith, 1970) and subsequent work on neutral networks (Schuster et al., 1994) showed that many genotypes can have similar fitness and that connectivity via neutral mutations provides a mechanism to perform a sustained exploration of

sequence space without fitness losses. In parallel, Eigen introduced the quasispecies framework (Eigen, 1971; Eigen et al., 1989), describing high-mutation-rate populations of replicators (Schuster and Sigmund, 1983; Schuster and Swetina, 1988), later applied to RNA viruses as clouds of related genotypes surrounding a master sequence (Domingo et al., 1978). In quasispecies populations, the effects of the underlying fitness landscape are attenuated by collective dynamics, allowing coexistence of mutants with diverse fitness values both in and out of mutation–selection equilibrium.

Efficient generation of genetic diversity is essential for long-term adaptability. Mechanisms that maintain extant diversity—such as high mutation rates or transposition bursts in response to unexpected selection pressures (Piacentini et al., 2014)—have likely been selected over evolutionary timescales (Ferrare and Good, 2024; Kumawat et al., 2025). In light of these observations, several key questions arise: To what extent are neutral spaces essential for evolvability? How constrained are populations to occupy fitness maxima? Are “valleys” in fitness landscapes true barriers to adaptation in high-dimensional, networked genotype spaces? And to what extent does finite population size limit the discovery of functional phenotypes?

2 Neutral Networks and Genotype-to-Phenotype Maps

2.1 *Synthetic Genotype-to-Phenotype Maps*

Maynard Smith’s idea of protein networks has been highly influential in subsequent characterizations and descriptions of both the structure of sequence spaces (induced by a variety of genotype-to-phenotype (GP) maps (Schuster et al. 1994; Lipman and Wilbur 1991; Arias et al. 2014; Greenbury et al. 2014; Fortuna et al. 2017) and the supposed constraints operating along evolutionary dynamics. In all analyses of GP maps there is the underlying assumption that mutations causing significant variation in the expressed phenotype will be rare or absent in populations, since they are deleterious and therefore selected against. Two ideas should be kept in mind in this context. First, natural selection acts on phenotypes, so significant evolutionary changes will only occur when phenotypes change. Since the fitness value of a phenotype is context-dependent, it is generally not possible to measure the value of a phenotype without testing its performance in specific environments. More often than not, even in constant environments, non-neutral mutations whose adaptive value is difficult to quantify are fixed in populations (Villanueva et al., 2022), occasionally modifying phenotypes in unforeseeable ways (Escarmís et al., 2008; Grant et al., 2021). Second, the presence of deleterious mutations in a population results from a competition between two temporal scales: the rate of generation of such mutations and the strength of purifying selection. This tension between the two scales is particularly relevant in ensembles of mutationally linked replicators, such as viruses, where the fast mutational scale translates into a persistent coexistence of variants bearing mutations with dissimilar effects on fitness (Lauring, 2020) and even different phenotypes (Duarte et al., 1994; Vignuzzi et al., 2006).

The previous observations notwithstanding, simple GP models have been tremendously enlightening for developing a more accurate picture of how genotypes map onto phenotypes. There are some properties shared by all navigable GP maps (Garcia-Martin et al., 2018), such as a negative correlation between genotypic evolvability and genotypic robustness, a positive correlation between phenotypic evolvability and phenotypic robustness, and a linear growth of phenotypic robustness with the logarithm of the number of genotypes expressing the same phenotype (what is usually referred to as “phenotype size”) (Manrubia et al., 2021). One of the most striking results of such models has been the quantification of phenotypic bias: Most of sequence space maps to a few highly abundant phenotypes, while an astronomical number of phenotypes are rare (few sequences map to them) and essentially invisible to evolution. These quantities are illustrated in Fig. 26.1 with the iconic example of the RNA sequence-to-secondary-structure GP map.

Phenotypic bias—the huge difference in phenotype sizes—hinted at by numerical analyses of GP maps, can be analytically derived in a general scenario where only one feature of the GP map matters: the fact that different sites in a sequence have different degrees of versatility (or neutrality) while yielding the same phenotype




Fig. 26.1 Neutral networks in the RNA sequence-to-secondary structure genotype–phenotype map. This grid provides a simplified representation of genotype networks, which in reality are far more complex and not planar. Each node represents a unique RNA sequence, and a single point mutation corresponds to a displacement to a neighboring node. All nodes of the same color (connected by thick black edges) represent RNA sequences that fold into the same secondary structure—i.e., different genotypes mapping to the same phenotype—and together constitute that phenotype’s *neutral network*. Three phenotypes (yellow, brown, and blue) occupy most of genotype space. The size of a neutral network ($N_1 = 15$ for the yellow folding, $N_2 = 10$ for the brown folding, and $N_3 = 12$ for the blue folding) is the number of sequences that produce the same phenotype. Many other, much smaller neutral networks (shown in various shades of gray) correspond to much rarer foldings, produced by only a few sequences (at most one or two in this example). Mutations within a neutral network allow divergence without losing a successful phenotype (hence without risking a fitness loss). Neutral networks usually traverse genotype space (e.g., the blue folding spans from top to bottom), and they get in touch with many other large phenotypes at one point or another (dashed red edges)

(Manrubia and Cuesta, 2017; Garcia-Martin et al., 2018; Martin and Ahnert, 2022; Catalán et al., 2023). The size of a phenotype can be approximated by the product of single-site versatility along the sequence, thus generating a log-normal distribution of phenotype sizes in almost any realistic GP map. Interestingly, various studies have shown that only large phenotypes (the rightmost tail of the log-normal distribution) are selected in natural systems (Jörg et al., 2008; Dingle et al., 2015). This highly relevant observation has two main consequences: The size of a phenotype must be incorporated as an important feature of its fitness, and the comparatively fewer large phenotypes are much more easily found through blind mutation than rare phenotypes—implying that highly frequent phenotypes dominate evolutionary dynamics (Cowperthwaite et al., 2008; Schaper and Louis, 2014; Catalán et al., 2020; Martin et al., 2024).

Neutral networks of large phenotypes have at least two additional evolutionary advantages: They are more robust on average, since the average number of neutral changes (the average degree of the neutral network $\langle k \rangle$) accepted by a sequence in a phenotype of size N is $\langle k \rangle \propto \log N$ (Aguirre et al., 2011; Catalán et al., 2023), and they also percolate the space of sequences, guaranteeing navigability and access to steadily growing novelty along evolution (Huynen et al., 1996; Huynen, 1996; Buchholz et al., 2017).

Still, what is the quantitative relevance of these potential advantages?

2.2 Phenotypic Robustness and Abundance

Let us begin by critically examining the effects of the increase in robustness with phenotype size, using the case of the RNA sequence-to-secondary structure map, where the relevant quantities have been analytically and numerically worked out. The goal of the following exercise is to estimate the fraction of single-point mutations that maintain the phenotype (thick black edges in Fig. 26.1) and the complementary fraction (i.e., those mutations causing a change in phenotype, light gray or dashed red edges in Fig. 26.1) for sequences of length L . This will quantitatively illustrate the difference between rare phenotypes (where most mutations are non-neutral, gray nodes in Fig. 26.1) and typically large phenotypes (colored nodes in Fig. 26.1), as found in natural systems.

On the basis of structural properties of RNA secondary structure, the expected number of neutral neighbors in a network of N genotypes mapping to the same secondary structure is, asymptotically in N (Catalán et al., 2023),

$$\langle k \rangle = c \log N + o((\log N)^{1/2}), \quad (26.1)$$

with $c \simeq 1.88$ (Aguirre et al., 2011). Bounds to minimum and maximum values of the coefficient c can be analytically obtained, yielding $c^{\min} = 1$ and $c^{\max} = 3/\log 4 \simeq 2.16$, respectively (Aguirre et al., 2011). In order to estimate the

average degree of frequent phenotypes, we need to estimate their average size N_{freq} as a function of sequence length. Analytical estimations show that the absolute number M_p of abundant phenotypes (those accumulating a fraction $p > 0.99$ of genotypes) grows as $M_p \simeq 1.33L^{-2}(1.68)^L$ (Catalán et al., 2023). A rough estimate of the typical size of those phenotypes is therefore given by

$$N_{\text{freq}} = \frac{4^L}{M_p} \simeq 0.75L^2(2.38)^L. \quad (26.2)$$

Neither M_p nor N_{freq} is significantly sensitive to the specific value of p , as long as $p \simeq 1$. Substituting the latter expression in Eq. (26.1), we obtain the expected neutrality for typical frequent phenotypes

$$\langle k \rangle \simeq 0.87cL + O(L^{1/2}) \simeq 1.64L + O(L^{1/2}). \quad (26.3)$$

Now, since every sequence of length L has $3L$ possible neighbors one mutation away, the expected fraction of neutral neighbors becomes $\rho_a = \langle k \rangle / (3L) \simeq 1.64/3 = 0.55$, with a correction $O(L^{-1/2})$. Asymptotically, therefore, we expect that around 45% of mutations yield new RNA secondary structures.

The latter figure compares well with that inferred from numerical studies, where the typical size of large, abundant phenotypes has been calculated (Dingle et al., 2015). Using the distribution of phenotype sizes corresponding to natural non-coding RNAs (Kin et al., 2007), the numerically estimated fraction ρ_n of neutral mutations (for sequence lengths $40 < L < 126$) reaches about $\rho_n \simeq 0.6$, so around 40% of mutations cause changes in phenotype.¹

Interestingly, both analytical ρ_a and numerical ρ_n estimations yield values independent of sequence length, revealing that robustness is not a quantity that can be tuned through insertions or deletions. In practice, the increase in neutrality of large phenotypes (with respect to small phenotypes, which are not observed in natural systems) implies that the mutation rate is slightly less effective in finding phenotypic novelty, requiring two or three attempts more than small phenotypes to escape the neutral network. It would be no surprise to learn that the mutation rate has been tuned along evolution to guarantee evolvability in large phenotypes—apparently the only ones occurring in natural populations.

¹Data in Fig. 3 of Dingle et al. (2015) yield a numerical estimate of the size N for each length L ; the average neutral degree of a phenotype of size N is then estimated as $\langle k \rangle_n = 1.88 \log N$, from which we get $\rho_n = \langle k \rangle_n / (3L)$; and ρ_n is a rough average over the values obtained for large L .

2.3 Percolation and Navigability in Neutral Networks

Multiple theoretical observations indicate that the neutral networks of large phenotypes contain sequences as distant as two random sequences could be. This property, together with the high dimensionality of sequence spaces (and the concomitant intertwining of phenotypes), conceptually guarantees that most (frequent) phenotypic alternatives are a few mutations away from at least some genomic region of the current phenotype. These properties result in network apposition—there are points of contact between any pair of phenotypes (Fontana, 2002) (dashed red edges in Fig. 26.1), a property empirically demonstrated (Schultes and Bartel, 2000)—and in the related property of shape space covering (novelty is a few mutations away) (Grüner et al., 1996a,b). What, then, is the role played by percolating phenotypes? Do populations need to explore neutral networks extensively in order to adapt successfully? This is not to suggest that neutral drift is irrelevant—on the contrary, it plays an important role (Bloom et al., 2007; Erdoğ an et al., 2024). Rather, the question concerns the extent of neutral drift required for adaptation and innovation. We illustrate this point below with a quantitative example.

Consider one of the shortest functional genomes, that of *Avocado sunblotch viroid* (ASBVd), with 246 nucleotides (Symons, 1981). The total number of different sequences of length $L = 246$ is $4^L \simeq 10^{148}$. A smaller figure corresponds to the number of sequences compatible with the RNA secondary structure of the ASBVd folded genome, with about 180 paired and 66 unpaired nucleotides, which is about 5×10^{41} (Garcia-Martin et al., 2018; Catalán et al., 2019): Still, the mass of all possible sequences in the ASBVd neutral network is about 10^4 times the total biomass on Earth. How many different sequences could have been produced since this viroid appeared? An upper bound can be estimated if we assume (i) ASBVd appeared simultaneously with land plants, some 470 My ago; (ii) the number of infected plants at any time is 10^{12} ,² (iii) the viroid load of any infected plant is 10^{12} ,³ and (iv) every replication of the genome includes on average one error (a new nucleotide). Assuming also that all generated sequences are different (though many of them might be identical or very similar and/or produce identical descent), we get an upper bound for the total number of explored sequences: 10^{36} . Though large, this number is still far from covering the complete neutral network of the ASBVd secondary structure, not to mention a significant fraction of the whole space

²Viroids infect a wide range of cultivated plants and ornamentals. The FAO estimates about 1.5×10^9 hectares of arable cropland; at a density of 10^5 – 10^6 plants per hectare, there are 10^{13} – 10^{14} plants in cultivation at any time. Assuming that 1–10% are infected at any time yields the figure of 10^{12} plants infected used in the order-of-magnitude estimation.

³Although precise quantification of viroid copy numbers per plant is seldom reported, high-throughput sequencing and biochemical assays reveal that systemic viroid infections can produce extremely high intracellular titers. For large hosts such as fruit trees, the product of cell number and per-cell viroid load readily exceeds 10^{12} genomes, supporting the use of that scale in upper bound estimates.

of constant L . Analogous figures for other genes and genomes yield values that make a systematic exploration of the set of sequences compatible with a phenotype absolutely out of reach (Louis, 2016).

What do these numbers mean? In light of the high diversity and plasticity of life and of the ability of organisms (notably microorganisms) to adapt to harsh and novel environments, it is clear that exploring most of sequence space is not a requisite for a population to thrive. As early studies with RNA postulated and illustrated numerically (Huynen et al., 1996) and more recent analyses suggest (Dryden et al., 2008), the specific identity of nucleotides (or amino acids) is not a requirement to attain function. Local exploration of variants—rather than exhaustive exploration of the full sequence space or even of an entire neutral network—must suffice.

3 Fitness Landscapes, Phenotype Accessibility, and High-Dimensional Constraints

A fitness landscape is a mapping from the genomic, multidimensional space to a real value representing fitness (Wright, 1931). Traditionally, fitness landscapes have been viewed as topographic representations where genotypes lie on the horizontal plane and fitness on the z -axis; this view has illustrated evolutionary dynamics for almost a century (Pigliucci, 2012). These landscapes portray the adaptive process as a parsimonious movement of populations that “traverse valleys,” “reach fitness peaks,” and adapt by incorporating mutations that pull them “uphill,” (see Fig. 26.2a). Inherent to these representations are concepts such as trapping in local maxima, peaks that can be attained, and, most of the time, fitness values with a sort of Platonic existence, independent of the state of the population. Although ample evidence shows that molecular evolution does not generally follow the iconic topographic fitness landscape, this metaphor remains deeply imprinted in our conceptual visualization of evolutionary dynamics.

3.1 Accessibility of Fitness Peaks in High Dimensions

Dynamics in multidimensional spaces differ fundamentally from two-dimensional projections. In high-dimensional genotype spaces, the accessibility of fitness peaks—i.e., the existence of monotonic mutational paths leading from arbitrary genotypes to the global fitness maximum—depends critically on the ruggedness and correlation structure of the landscape. Analytical and numerical studies of random fitness models have shown that in uncorrelated or weakly correlated landscapes, such as the House-of-Cards or NK models with fixed local interactions, the probability that any strictly increasing path reaches the global optimum decays exponentially with sequence length L (Franke et al., 2011; Krug, 2019; Schmiegel

and Krug, 2023). This loss of accessibility occurs because each beneficial step must satisfy a sequence of rare inequalities in fitness orderings, and in high dimensions these constraints multiply, making global fitness peaks effectively unreachable by purely uphill walks. Increasing the alphabet size K (allowing more allelic states per site) or introducing a global additive gradient (as in the Mount Fuji model) can substantially increase the number of accessible trajectories, yet these conditions are atypical of natural, epistatic landscapes (Zagorski et al., 2016; Franke et al., 2011).

While small empirical landscapes often exhibit multiple accessible routes (Szendro et al., 2013), the generic expectation for large genomes is that strictly monotonic adaptive paths are exceedingly rare. Reaching the global maximum typically requires transient declines in fitness or exploration of neutral or nearly neutral networks. Hence, the intuitive Wrightian picture of populations climbing continuous mountains is misleading in high-dimensional sequence spaces: Despite the astronomical number of possible mutational paths, most are blocked by sign epistasis, and purely uphill trajectories almost never connect distant genotypes (Krug, 2019).

3.2 *Invisible Peaks and Entropic Barriers*

A more accurate description of evolutionary dynamics on sequence spaces should combine rough fitness landscapes with the partition of the space induced by mapping genotypes to phenotypes. Although genotype-to-fitness landscapes are often the default representation, considering the intermediate mapping to phenotypes illustrates an additional potential difficulty in reaching high-fitness solutions. The unified picture is that of adaptive multiscales (Catalán et al., 2017), where each genotype maps to one phenotype (in a simplified many-to-one GP map), and all genotypes belonging to the same phenotype take the same fitness value in a specific environment. In this scenario, the highly skewed distribution of phenotype sizes comes into play: Only the largest phenotypes are relevant to evolution; fitter phenotypes, if sufficiently smaller, are almost never found. In other words, even if the global fitness peak (understood as a degenerate “peak” corresponding to all genotypes mapping to that phenotype) is accessible, it is likely lost in a sea of astronomically more abundant alternatives.

3.3 *Phenotypic Accessibility and the Arrival of the Frequent*

Mean-field models of GP maps provide a remarkably compact description of the high-dimensional dynamics of molecular evolution. In such models, the complex local structure of neutral networks is averaged out, and the transition probability $\phi_{\xi \rightarrow \chi}$ from phenotype ξ to phenotype χ upon a point mutation can be computed as an

average over all genotypes realizing ξ (Schaper and Louis, 2014; Greenbury et al., 2016; Manrubia and Cuesta, 2015; Huynen et al., 1996). This approach predicts that, to a first approximation,

$$\phi_{\xi \rightarrow \chi} \propto f_{\chi}, \quad (26.4)$$

where f_{χ} is the global frequency (or relative neutral set size) of phenotype χ , i.e., the fraction of sequence space that maps to it. Because the distribution of f_{χ} values typically spans many orders of magnitude and is approximately log-normally distributed in RNA and other GP maps, the corresponding transition probabilities vary comparably. As a result, most of the mutational flux from any given phenotype leads to a very small subset of abundant, highly connected phenotypes, while rare phenotypes—those with small size N —are practically inaccessible, even when they would confer large fitness advantages. This bias, termed the *arrival of the frequent* (Schaper and Louis, 2014), implies that the spectrum of available variation is determined primarily by the statistical geometry of the GP map rather than by selective gradients on fitness. Note that the mean-field estimation is remarkably optimistic regarding transition probabilities: If there is a small number of links between the two phenotypes, and considering the huge sizes of the latter (even small ones), populations may evolve for extremely long times away from the apposition regions where the two phenotypes are one mutation apart.

A simple numerical example illustrates the magnitude of this bias in the mean-field scenario. Consider sequences of length $L = 100$ with a four-letter alphabet (as in RNA), yielding $4^{100} \simeq 10^{60}$ possible genotypes. A typical large phenotype has size $N_{\text{freq}} \simeq 10^{42}$, the most abundant phenotype size is 10^{32} , and typical small phenotypes have size $N_{\text{rare}} \simeq 10^{22}$ (Dingle et al., 2015). For a rare phenotype with a neutral set size 10^{20} times smaller, both theory and simulation suggest that, to a first approximation, the probability of a mutation leading into this phenotype scales roughly with its relative size N , so $\phi_{\text{rare}}/\phi_{\text{freq}} \simeq 10^{-20}$.⁴ Hence, even if a population explores 10^{12} *different* genotypes through mutation and drift, the expected number of encounters with that rare phenotype remains negligibly small. Conversely, neutral exploration within frequent phenotypes dominates: The population will perform an effectively *local* random walk within large, percolating neutral networks that connect vast regions of sequence space, even if large excursions from the initial sequence are uncommon.

Analytical results allow estimating the ratio between the sizes of frequent and rare phenotypes as the sequence length increases. For RNA, the distribution of the natural logarithm of phenotype sizes $\log N$ follows a normal distribution with mean $\mu_L = \mu_1 L + O(1)$ and standard deviation $\sigma_L = \sigma_1 L^{1/2} + O(L^{-1/2})$, where $\mu_1 = 0.2865$ and $\sigma_1 = 0.4434$ (Garcia-Martin et al., 2018). Then, the typical size of large phenotypes can be shown to be (Catalán et al., 2023)

⁴This ratio could be corrected by the fraction of neutral mutations for each phenotype size, but the correction is just a factor of 2, thus not modifying the order of magnitude.

$$\log N_{\text{freq}} \simeq \mu_L + \sigma_L^2 - z\sqrt{2}\sigma_L, \quad (26.5)$$

with $z \in (1.5, 2.7)$ for the fraction of genotypes included in the set of frequent phenotypes, $p \in (0.99, 0.9999)$. By symmetry, since the distribution of phenotype sizes peaks at μ_L , the typical size of rare phenotypes can be analogously estimated as

$$\log N_{\text{rare}} \simeq \mu_L - \sigma_L^2 + z\sqrt{2}\sigma_L. \quad (26.6)$$

Hence, the mean-field estimate of the ratio of rare to frequent phenotype sizes is

$$\frac{\phi_{\text{rare}}}{\phi_{\text{freq}}} \simeq \exp\left(-2\sigma_1^2 L + 2\sqrt{2}z\sigma_1\sqrt{L}\right). \quad (26.7)$$

This estimation yields larger relative values for $L = 100$ than the rough numerical estimate obtained above, since it estimates smaller and larger phenotype sizes for frequent and rare phenotypes, respectively. However, the important result is that the ratio decreases exponentially with L , broadening the separation between the size of frequent and rare phenotypes and making rare phenotypes increasingly difficult to meet. For $L = 246$, as in the Avocado sunblotch viroid example, and taking $z = 2$, $\phi_{\text{rare}}/\phi_{\text{freq}} \simeq 4 \times 10^{-34}$.

In summary, small, isolated phenotypes remain essentially invisible. Canalization or selective bias could, in principle, guide trajectories toward fitter phenotypes, but only within the subset of variation that the GP map makes accessible. This subset is essentially covered by other frequent phenotypes.

Even if a population were artificially placed on one of these small, high-fitness phenotypes—analogueous to fixing an adaptive RNA or protein structure in the laboratory—the phenotype would be rapidly lost under mutation pressure. Because most mutational trajectories leave the small neutral set, the population diffuses away in a few generations toward more abundant phenotypes with larger mutational robustness. This loss is fundamentally entropic: In the same way that systems in statistical physics evolve toward macrostates with more microstates, evolutionary dynamics naturally drift toward regions of sequence space with greater combinatorial weight. Thus, both the difficulty of locating rare phenotypes and the instability of remaining on them ensure that the “arrival of the frequent” dominates evolutionary exploration, limiting accessibility to the global fitness maximum (which, by probability, corresponds to the set of non-frequent phenotypes) except under extreme, nongeneric conditions—an outcome consistent with results from models of fitness landscapes in high dimensions, where ruggedness and dimensionality conspire to render global peaks exponentially rare and typically unreachable (Franke et al., 2011; Nowak and Krug, 2013).

4 Efficient Adaptation Within High-Abundance Phenotypes

The entropic trapping effect caused by the large size of abundant phenotypes was described long ago (Schuster and Swetina, 1988) and observed in multiple numerical studies (Wolff and Krug, 2009; Khatri et al., 2009; Wilke et al., 2001; Codoñer et al., 2006; Cowperthwaite and Meyers, 2007; Schaper and Louis, 2014; Catalán et al., 2020), where this effect manifests in various forms. Terms such as *phenotypic entropy*, *redundancy*, or *landscape flatness* have been used to conceptually describe the number of genotypes that realize a given phenotype and, consequently, its mutational robustness. A pioneering theoretical analysis (Schuster and Swetina, 1988) considered the competition between two quasispecies corresponding to distinct phenotypes and showed that the phenotype with the highest replication rate could nevertheless be outcompeted by a slightly less fit one if the latter occupied a larger and more connected neutral network. Too low robustness, they demonstrated, could be fatal at high mutation rates. Note that in this analysis, the phenotype size was not included as a feature of fitness, and only the replicative ability of the genotypes within the phenotype was considered. An analytical treatment (Wolff and Krug, 2009) later quantified the conditions for this crossover between fitness and robustness, showing that increasing mutation rates can drive the dominance of phenotypes with lower intrinsic fitness but broader mutational neighborhoods and that this transition is closely related to the onset of the error catastrophe.

Other studies have examined analogous effects in different biological and computational contexts. In models of spatial gene regulation, it was shown that finite populations often fail to converge to the globally optimal phenotype when multiple equivalent solutions exist, an outcome interpreted as entropic and arising from the multiplicity of equally viable configurations (Khatri et al., 2009). The same principle underlies the so-called survival of the flattest phenomenon (Wilke et al., 2001), where, at sufficiently high mutation rates, populations evolve toward flatter regions of the fitness landscape, prioritizing mutational robustness over maximal replication speed. This prediction was later confirmed experimentally with RNA viroids (Codoñer et al., 2006), where lineages displaying slower replication but higher neutrality outcompeted faster yet fragile variants, showing that fitness can be maximized by minimizing the mutational load rather than by maximizing growth rate.

Computational analyses of mutation–selection dynamics have further clarified the evolutionary implications of phenotype size. It has been shown that adaptation tends to proceed through abundant phenotypes—a process also termed the “ascent of the abundant” (Cowperthwaite and Meyers, 2007)—and that phenotypic innovation is biased toward more frequent phenotypes in genotype–phenotype maps, the effect already described as the “arrival of the frequent” (Schaper and Louis, 2014). More recent analytical work (Catalán et al., 2020) has demonstrated that this bias persists across a broad range of mutation–selection regimes, confirming that phenotypic accessibility, rather than differences in maximal fitness, largely determines the evolutionary trajectories available to finite populations.

In summary, only a vanishingly small fraction of highly abundant phenotypes is explored along evolution and thus presented to natural selection. The question of whether this restriction limits, and to what extent, the localization of phenotypes with sufficiently high fitness was addressed in a study where the relative weight of phenotype size versus replicative ability was analytically and numerically explored (Catalán et al., 2023). Returning to the case of RNA sequence-to-secondary structure GP maps, the fraction u of phenotypes visible to evolution diminishes exponentially with L as

$$u \simeq \frac{0.9}{L^{1/2}} (1.1)^{-L}. \quad (26.8)$$

However, since the total number of different open secondary structures for sequences of length L grows as $1.48L^{-3/2}(1.85)^L$, the absolute number of different, large phenotypes covered by a fraction $p \simeq 1$ of genotypes grows exponentially with L as $1.33L^{-2}(1.68)^L$.

Assuming, as in Catalán et al. (2023), that replicative ability values follow a normal distribution with average $\langle r \rangle$ and variance σ_r^2 , even in the restricted set of large phenotypes visible to natural selection, the maximum value of r that can be attained is only about 10% lower than in the whole ensemble of phenotypes. In other words, limiting the evolutionary search to large phenotypes does not entail a large fitness loss, while it can yield a significant gain in robustness and navigability.

The exploration of phenotypic diversity is not only restricted to large phenotypes, but it is also local in sequence space. No better demonstration of the locality of evolution exists than gene phylogenies, where common ancestors occurred billions of years ago. Retaining memory of this distant past indicates that extant, successful functions were not far from the sequences where the search started. This observation relates to the idea of shape space covering discussed above, which appears to be a general feature of the broad sequence-to-function map: New functions may lie a few mutations away from current genotypes in high-dimensional spaces. Support for this expectation, originally derived from simple GP maps, has come from experimental results showing that functional genes can be rapidly generated from nongenic sequences (Carvunis et al., 2012; Vakirlis et al., 2020), that transposable elements can be domesticated to perform specific functions in their host (Sinzelle et al., 2009), or that promoters can emerge from random sequences (Yona et al., 2018), even in the absence of sequence diversity, simply through successive cycles of mutation, enrichment, and selection (Wachowius et al., 2021). Thus, high-dimensional sequence spaces allow the discovery and fixation of molecular function diversity observed in current organisms, speeding up adaptive rates so much as to allow organisms to thrive in steadily changing environments.

5 Hierarchical Genotype Networks in Empirical Populations

A conceptual metaphor that helps us understand how evolutionary dynamics unfold in sequence spaces needs to integrate the main elements discussed in previous sections: (i) Sequence spaces are so vast that only a minute fraction is explored by natural populations, (ii) neutral networks are astronomically large, and natural populations only visit a local, tiny fraction of the whole, (iii) spaces of high dimensionality are so deeply interwoven that most (large) phenotypes can be found a few mutations away, (iv) as sequence length increases, high-fitness maxima become increasingly difficult to attain, (v) the fraction of (large) phenotypes explored by genotypes of length L decreases exponentially with L , though at the same time, the absolute number of different phenotypes increases exponentially, and (vi) the latter restriction has a small impact on the replicative ability values that populations can achieve. Establishing a mindset within this scenario should also be helpful when interpreting empirical results and should prove especially relevant when heterogeneity is intrinsic to the population, such as in viruses and other organisms with high genotypic diversity. These populations should comply with the restrictions above while managing to explore sufficient evolutionary innovation or variants that allow their adaptation to changing conditions. A pictorial summary of how the elements above modify suitable representations of quasispecies evolution on sequence spaces is represented in Fig. 26.2.

Advances in sequencing technologies have boosted the detailed description of genetic diversity in viral populations in particular. This has allowed the characterization, to varying but deep detail, of the within-host structure of viral populations (Baaijens et al., 2017; Zanini et al., 2017; Delgado et al., 2021) and their phylodynamics (Attwood et al., 2022; Sun et al., 2025). Recently, deep sequencing of the $Q\beta$ bacteriophage has revealed a hierarchical structure with implications for viral evolutionary dynamics and, in particular, for the search process in sequence spaces (Seoane et al., 2026).

$Q\beta$ populations exhibit robust organizational patterns along evolution, even when new mutations with a supposed adaptive value become fixed. Figure 26.3 represents a snapshot of a genotype network of a deep-sequenced $Q\beta$ population. At any time point in a process of adaptation to new environments, genomes cluster around the most abundant sequence, which appears to continuously regenerate diversity, acting as a diversity source and populating its neighborhood at increasing mutational distance in a process dominated by simple diffusion: Sequence abundance decays exponentially with Hamming distance to the most abundant sequence. A remarkable observation is that, on average, no significant correlation between the abundance of variants and the type of mutation they carry (synonymous or non-synonymous) is found. In this out-of-equilibrium state, the search process is dominated by reaction (the most abundant sequence produces a larger number of offspring) and diffusion (through mutation), with a minor effect of the underlying fitness landscape. As new mutations fix in the population, the whole hierarchical structure is rapidly

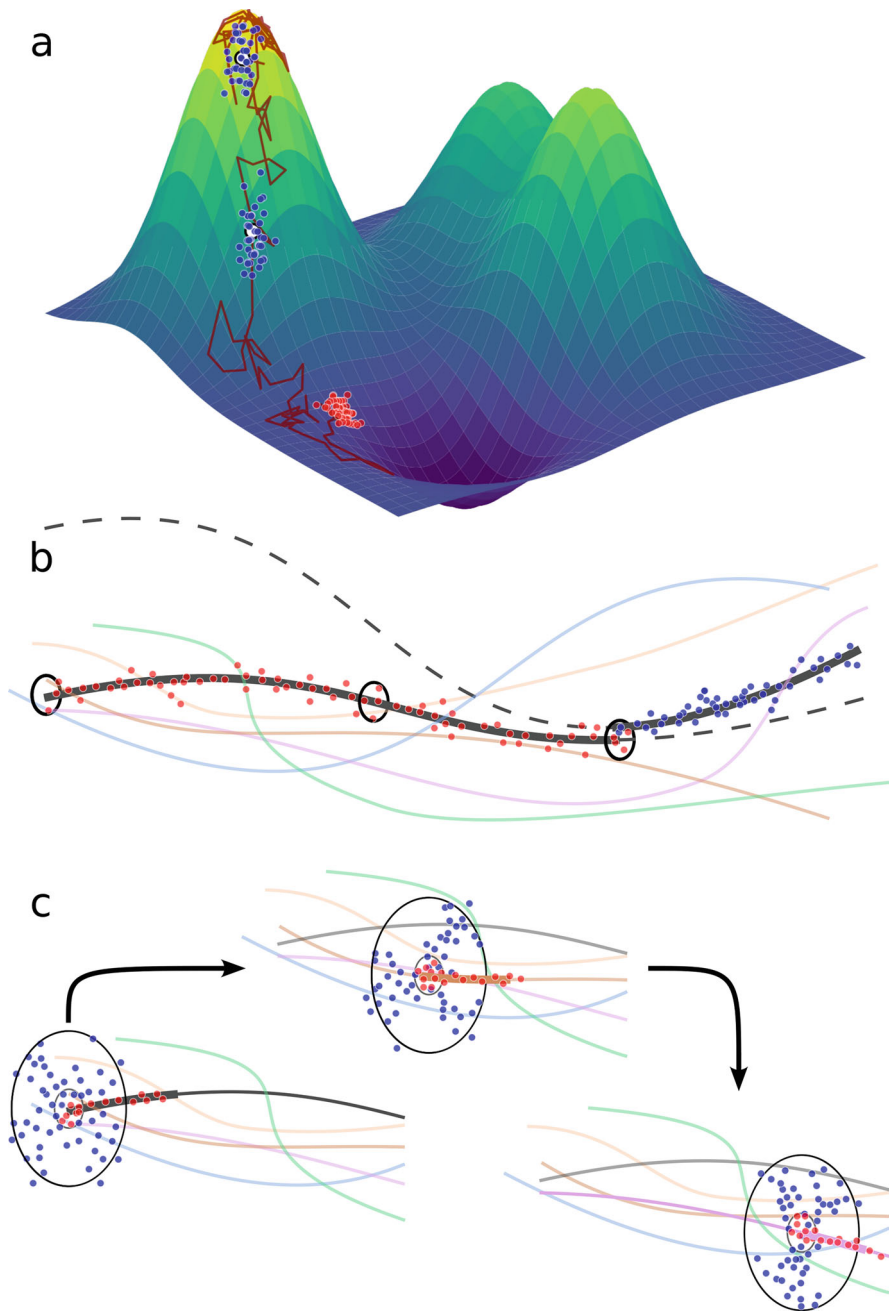


Fig. 26.2 Historical illustration of evolving viral quasispecies. Our understanding of how viral quasispecies explore sequence space has evolved over time as empirical evidence has accumulated. (a) In early fitness-landscape models, evolutionary dynamics were primarily described in terms of fitness increases. A simulated adaptive walk (brown curve with sampled red and blue points) illustrates this view. A viral population is assumed to start at low fitness, with individual variants independently sampling the local landscape. Fitter variants are selectively amplified, allowing the

reconstructed around the new center. Figure 26.3 illustrates the uncommon instance when the population transitions between two dominant haplotypes.

These observations illustrate how heterogeneous populations explore sequence space and, with it, new potential phenotypes: Little use of neutral paths seems to occur, with sequential displacements being apparently more dependent on mutations with at least some adaptive advantage. At the same time, almost all possible variants one- and two-mutations away, and a significant fraction at larger distances, are generated by a massive production of mutants before purifying selection can eliminate low-fitness variants. Therefore, fitness “valleys” no longer confine the population to high-fitness variants. In this scenario, the concept of a fitness peak becomes diffuse: Population dynamics only see a coarse representation of the landscape, where the fitness of individual sequences is no longer the unique determinant of their fate—whether maintained or discarded. This provides an alternative view of quasispecies as units of selection, since their structure and dynamics are in mutual feedback to guarantee persistence.

6 Conclusions

Evolution and adaptation occur primarily at a local level in sequence space. The astronomical size of sequence spaces, even of the subsets of genotypes mapping to the same phenotype, strongly suggests that this cannot be otherwise. Evidence for limited displacements in these vast spaces since functional solutions were first located is provided by gene phylogenies that can be traced hundreds of millions of years back in time: Long excursions through neutral drift are not essential for adaptation.

Fig. 26.2 (continued) population as a whole to climb toward higher fitness peaks. Stochastic effects may occasionally enable escape from local optima. **(b)** When most mutations are nonviable or incur large fitness losses, the simple fitness-ascent picture becomes problematic. This motivated a shift in emphasis from fitness progression to neutral navigation. Neutral networks (colored curves) span genotype space, connecting widely different sequences that produce the same phenotype through successive neutral mutations. A quasispecies can drift along such a network (thick black curve and red dots) without loss of fitness, while continually sampling nearby phenotypes. Occasionally, a previously distant neutral network (dashed black curve) comes into close proximity, offering access to higher fitness. The population then shifts to and explores this new network (blue dots). **(c)** Recent empirical evidence suggests that viral quasispecies do not primarily rely on neutral network connectivity to navigate fitness landscapes. Instead, they generate a large fraction of all possible variants within a relatively broad mutational neighborhood (outer black circles) around the most abundant sequence, which is itself highly populated (blue dots). Neutral variants (red dots) are explored within the same radius, rather than preferentially. The schematic separation between neutral and non-neutral variation is shown for clarity; in reality, neutral networks are interwoven with non-neutral neighborhoods. Shifts in the dominant (reference) sequence occur frequently (twice in this illustration), after which the statistical structure of the explored neighborhood is rapidly reestablished




Fig. 26.3 An empirical genotype network. Each node in this network represents a unique sequence (haplotype) of a $Q\beta$ phage population; two haplotypes are connected if they differ in just one nucleotide. Data correspond to passage 60 of a population adapting to a temperature of 43 °C. A full description of the deep sequencing data and the curation protocol can be found in Seoane et al. (2026). The node size is proportional to the logarithm of the abundance of the haplotype (the number of sequences identical to that haplotype), and color conveys the same information. The spring layout (which positions nodes by treating them as masses connected by springs and minimizing tensions) suggests community structure around a series of salient nodes. These seem to act as *sources* of viral replicas which, through mutation, populate a neighborhood around each network hub. The self-similar structure of the network is reflected, among others, in a robust power-law degree distribution (Seoane et al., 2026)

Despite this local exploration, the high dimensionality of sequence spaces causes abundant phenotypes to be highly interwoven, with regions of contact that are frequent and occur in many separated regions of sequence space. Consequently, a limited exploration of genotype space does not imply a limited exploration of phenotype space. This conclusion extends to the exploration of fitness space,

assuming that the distribution of fitness values has a well-defined average (Catalán et al., 2023)—an assumption supported by the variations in functionality observed in natural systems.

A particularly striking result is that Maynard Smith's notion of a connected protein space, which allows the sustained accumulation of mutations without substantial fitness loss, is not required for the evolution and adaptation of viral populations. Instead, it suffices that phenotype abundances be highly biased, ensuring that frequent phenotypes lie close to one another in sequence space. Viral populations generate and explore all available diversity at short mutational distances from fit, abundant sequences, largely independently of the detailed structure of the underlying fitness landscape. This finding is relevant for understanding the emergence of function in populations of replicators, both within cellular environments (as may have occurred during the emergence of viroids (Catalán et al. 2019)) and in pre-cellular evolutionary contexts (Manrubia, 2022).

We may even entertain the idea that percolation of sequence space by neutral networks and the notion that it facilitates navigation of genotype spaces are merely spandrels of the properties of large phenotypes that are actually relevant for evolution: They exist due to their abundance and persist due to entropic principles. Only a minute fraction of neutral networks is ever visited in evolutionary time, a fact that does not contradict the emergence of molecular function through repeated selection cycles on random sequences: Functional sequences are often just a few mutations away.

Acknowledgements The authors acknowledge the contributions of Jacobo Aguirre, Iker Atienza, Pablo Catalán, Juan Antonio García-Martín, Ester Lázaro, Samuel Martínez-Alcalá, Henry Secaira-Morcho, and Ariadna Villanueva to the development of the unified picture presented here. This work was funded by MICIU/AEI/10.13039/501100011033 and by ERDF/EU (PGE) “A way of making Europe,” through grants PID2023-147963NB-C21 (SM), PID2023-153225NA-I00 (LFS), and PID2022-141802NB-I00 (BASIC) (JAC).

References

- Agashe D, Sane M, Singhal S (2023) Revisiting the role of genetic variation in adaptation. *Am Natural* 202(4):486–502. <https://doi.org/10.1086/726012>
- Aguirre J, Buldú JM, Stich M, Manrubia SC (2011) Topological structure of the space of phenotypes: the case of RNA neutral networks. *PLoS ONE* 6:e26324
- Arias CF, Catalán P, Manrubia S, Cuesta JA (2014) toyLIFE: a computational framework to study the multi-level organisation of the genotype-phenotype map. *Sci Rep* 4:7549
- Attwood SW, Hill SC, Aanensen DM, Connor TR, Pybus OG (2022) Phylogenetic and phylodynamic approaches to understanding and combating the early SARS-CoV-2 pandemic. *Nat Rev Genet* 23(9):547–562. <https://doi.org/10.1038/s41576-022-00483-8>
- Baaijens JA, Aabidine AZE, Rivals E, Schönhuth A (2017) De novo assembly of viral quasispecies using overlap graphs. *Genome Res* 27(5):835–848
- Bloom JD, Romero PA, Lu Z, Arnold FH (2007) Neutral genetic drift can alter promiscuous protein functions, potentially aiding functional evolution. *Biol Dir* 2:17

- Bordería AV, Stapleford KA, Vignuzzi M (2011) RNA virus population diversity: implications for inter-species transmission. *Current Opin Virol* 1(6):643–648. <https://doi.org/10.1016/j.coviro.2011.09.012>. <https://www.sciencedirect.com/science/article/pii/S1879625711001192>. Innate immunity/Antivirals and resistance/Emerging viruses
- Braun T, Bordería AV, Barbezange C, Vignuzzi M, Louzoun Y (2018) Long-term context-dependent genetic adaptation of the viral genetic cloud. *Bioinformatics* 35(11):1907–1915 <https://doi.org/10.1093/bioinformatics/bty891>
- Buchholz PCF, Fademrecht S, Pleiss J (2017) Percolation in protein sequence space. *PLOS ONE* 12(12):1–11 <https://doi.org/10.1371/journal.pone.0189646>
- Carvunis AR, Rolland T, Wapinski I, Calderwood MA, Yildirim MA, Simonis N, Charlotteaux B, Hidalgo CA, Barbette J, Santhanam B, Brar GA, Weissman JS, Regev A, Thierry-Mieg N, Cusick ME, Vidal M (2012) Proto-genes and de novo gene birth. *Nature* 487(7407):370–374
- Catalán P, Arias CF, Cuesta JA, Manrubia S (2017) Adaptive multiscales: an up-to-date metaphor to visualize molecular adaptation. *Biol Direct* 12:7
- Catalán P, Elena SF, Cuesta JA, Manrubia S (2019) Parsimonious scenario for the emergence of viroid-like replicons de novo. *Viruses* 11:425
- Catalán P, Manrubia S, Cuesta JA (2020) Populations of genetic circuits are unable to find the fittest solution in a multilevel genotype-phenotype map. *J R Soc Interface* 17:20190843
- Catalán P, Garcia-Martin JA, Aguirre J, Cuesta JA, Manrubia S (2023) Entropic contribution to phenotype fitness. *J Phys A Math Theoret* 56(34):345601. <https://doi.org/10.1088/1751-8121/ace8d6>
- Codoñer FM, Darós JA, Solé RV, Elena SF (2006) The fittest versus the flattest: experimental confirmation of the quasispecies effect with subviral pathogens. *PLoS Pathog* 2:e136
- Cowperthwaite MC, Meyers LA (2007) How mutational networks shape evolution: lessons from RNA. *Annu Rev Ecol Syst* 38:203–230
- Cowperthwaite MC, Economo EP, Harcombe WR, Miller EL, Meyers LA (2008) The ascent of the abundant: how mutational networks constrain evolution. *PLoS Comput Biol* 4:e1000110
- Cvijović I, Nguyen Ba AN, Desai MM (2018) Experimental studies of evolutionary dynamics in microbes. *Trends Genet* 34(9):693–703. <https://doi.org/10.1016/j.tig.2018.06.004>
- Delgado S, Perales C, García-Crespo C, Soria ME, Gallego I, de Ávila AI, Martínez-González B, Vázquez-Sirvent L, López-Galíndez C, Morán F, Domingo E (2021) A two-level, intramutant spectrum haplotype profile of hepatitis C virus revealed by self-organized maps. *Microbiol Spectrum* 9(3):e01459–21. <https://doi.org/10.1128/Spectrum.01459-21>. <https://journals.asm.org/doi/abs/10.1128/spectrum.01459-21>
- Dingle K, Schaper S, Louis AA (2015) The structure of the genotype-phenotype map strongly constrains the evolution of non-coding RNA. *Interface Focus* 5:20150053
- Domingo E (ed) (2006) *Quasispecies: concept and implications for virology*. Springer, Berlin
- Domingo E, Sabo D, Taniguchi T, Weissmann C (1978) Nucleotide sequence heterogeneity of an RNA phage population. *Cell* 13(4):735–744. [https://doi.org/10.1016/0092-8674\(78\)90223-4](https://doi.org/10.1016/0092-8674(78)90223-4)
- Dryden DT, Thomson AR, White JH (2008) How much of protein sequence space has been explored by life on earth? *J R Soc Interface* 5(25):953–956. <https://doi.org/10.1098/rsif.2008.0085>. <https://royalsocietypublishing.org/doi/abs/10.1098/rsif.2008.0085>. <https://royalsocietypublishing.org/doi/pdf/10.1098/rsif.2008.0085>
- Duarte EA, Novella IS, Ledesma S, Clarke DK, Moya A, Elena SF, Domingo E, Holland JJ (1994) Subclonal components of consensus fitness in an RNA virus clone. *J Virol* 68(7):4295–4301. <https://doi.org/10.1128/jvi.68.7.4295-4301.1994>. <https://journals.asm.org/doi/abs/10.1128/jvi.68.7.4295-4301.1994>. <https://journals.asm.org/doi/pdf/10.1128/jvi.68.7.4295-4301.1994>
- Eigen M (1971) Selforganization of matter and the evolution of biological macromolecules. *Naturwissenschaften* 58:465–523
- Eigen M, McCaskill J, Schuster P (1989) The molecular quasi-species. *Adv Chem Phys* 75:149–263
- Erdoğan AN, Dasmeh P, Socha RD, Chen JZ, Life BE, Jun R, Kiritchkov L, Kehila D, Serohijos AWR, Tokuriki N (2024) Neutral drift upon threshold-like selection promotes variation in

- antibiotic resistance phenotype. *Nature Communications* 15(1):10813. <https://doi.org/10.1038/s41467-024-55012-4>
- Escarmís C, Lázaro E, Arias A, Domingo E (2008) Repeated bottleneck transfers can lead to non-cytocidal forms of a cytopathic virus: implications for viral extinction. *J Mol Biol* 376(2): 367–379. <https://doi.org/10.1016/j.jmb.2007.11.042>. <https://www.sciencedirect.com/science/article/pii/S0022283607015136>
- Ferrare JT, Good BH (2024) Evolution of evolvability in rapidly adapting populations. *Nat Ecol Evol* 8(11):2085–2096. <https://doi.org/10.1038/s41559-024-02527-0>
- Fontana W (2002) Modelling ‘evo-devo’ with RNA. *BioEssays* 24:1164–1177
- Fortuna MA, Zaman L, Ofria C, Wagner A (2017) The genotype-phenotype map of an evolving digital organism. *PLoS Comput Biol* 13(2):e1005414
- Franke J, Klozer A, de Visser JAGM, Krug J (2011) Evolutionary accessibility of mutational pathways. *PLoS Comput Biol* 7(8):e1002134. <https://doi.org/10.1371/journal.pcbi.1002134>
- García-Martin JA, Catalán P, Cuesta JA, Manrubia S (2018) Statistical theory of phenotype abundance distributions: a test through exact enumeration of genotype spaces. *Europhys Lett* 123:28001
- Grant NA, Magid AA, Franklin J, Dufour Y, Lenski RE (2021) Changes in cell size and shape during 50,000 generations of experimental evolution with *Escherichia coli*. *J Bacteriol* 203(10): 10.1128/jb.00469–20. <https://doi.org/10.1128/jb.00469-20>. <https://journals.asm.org/doi/abs/10.1128/jb.00469-20>, <https://journals.asm.org/doi/pdf/10.1128/jb.00469-20>
- Greenbury SF, Johnston IG, Louis AA, Ahnert SE (2014) A tractable genotype-phenotype map for the self-assembly of protein quaternary structure. *J R Soc Interface* 11:20140249
- Greenbury SF, Schaper S, Ahnert SE, Louis AA (2016) Genetic correlations greatly increase mutational robustness and can both reduce and enhance evolvability. *PLoS Comput Biol* 12: e1004773
- Grüner W, Giegerich R, Strothmann D, Reidys C, Weber J, Hofacker IL, Stadler PF, Schuster P (1996a) Analysis of RNA sequence structure maps by exhaustive enumeration. I. Neutral networks. *Monatsh Chem* 127:355–374
- Grüner W, Giegerich R, Strothmann D, Reidys C, Weber J, Hofacker IL, Stadler PF, Schuster P (1996b) Analysis of RNA sequence structure maps by exhaustive enumeration II. Structures of neutral networks and shape space covering. *Monatsh Chem* 127:375–389
- Gu H, Quadeer AA, Krishnan P, Ng DYM, Chang LDJ, Liu GYZ, Cheng SMS, Lam TTY, Peiris M, McKay MR, Poon LLM (2023) Within-host genetic diversity of SARS-CoV-2 lineages in unvaccinated and vaccinated individuals. *Nat Commun* 14(1):1793. <https://doi.org/10.1038/s41467-023-37468-y>
- Harris KB, Flynn KM, Cooper VS (2021) Polygenic adaptation and clonal interference enable sustained diversity in experimental *Pseudomonas aeruginosa* populations. *Mol Biol Evol* 38(12):5359–5375. <https://doi.org/10.1093/molbev/msab248>. <https://academic.oup.com/mbe/article-pdf/38/12/5359/41714606/msab248.pdf>
- Huynen MA (1996) Exploring phenotype space through neutral evolution. *J Mol Evol* 43:165–169
- Huynen MA, Stadler PF, Fontana W (1996) Smoothness within ruggedness: the role of neutrality in adaptation. *Proc Natl Acad Sci USA* 93:397–401
- Jones LR (2025) Intra-host variability of SARS-CoV-2: patterns, causes and impact on COVID-19. *Virology* 603:110366. <https://doi.org/10.1016/j.virol.2024.110366>. <https://www.sciencedirect.com/science/article/pii/S0042682224003908>
- Jörg T, Martin OC, Wagner A (2008) Neutral network sizes of biological RNA molecules can be computed and are not atypically small. *BMC Bioinf* 9:464
- Khatri BS, McLeish TCB, Sear RP (2009) Statistical mechanics of convergent evolution in spatial patterning. *Proc Natl Acad Sci USA* 106(24):9564–9569
- Kim N, Ma J, Kim W, Kim J, Belenky P, Lee I (2024) Genome-resolved metagenomics: a game changer for microbiome medicine. *Exp Mol Med* 56(7):1501–1512. <https://doi.org/10.1038/s12276-024-01262-7>
- Kimura M (1968) Evolutionary rate at the molecular level. *Nature* 217:624–626

- Kin T, Yamada K, Terai G, Okida H, Yoshinari Y, Ono Y, Kojima A, Kimura Y, Komori T, Asai K (2007) fRNADb: a platform for mining/annotating functional RNA candidates from non-coding RNA sequences. *Nucleic Acids Res* 35(suppl 1):D145–D148
- Krug J (2019) Accessibility percolation in random fitness landscapes. In: *Statistical physics and evolutionary biology*. Springer, Berlin. <https://arxiv.org/abs/1903.11913>. arXiv:1903.11913
- Kumawat B, Lalejini A, Acosta MM, Zaman L (2025) Evolution takes multiple paths to evolvability when facing environmental change. *Proc Natl Acad Sci* 122(1):e2413930121. <https://doi.org/10.1073/pnas.2413930121>. <https://www.pnas.org/doi/abs/10.1073/pnas.2413930121>. <https://www.pnas.org/doi/pdf/10.1073/pnas.2413930121>
- Laguna-Castro M, Somovilla P, López-Muñoz V, Pacios LF, Lázaro E (2025) Navigating the fitness landscape: host density, epistasis, and clonal interference drive divergent evolutionary pathways in phage Q β . *Int J Mol Sci* 26(18). <https://doi.org/10.3390/ijms26189020>. <https://www.mdpi.com/1422-0067/26/18/9020>
- Lang GI, Rice DP, Hickman MJ, Sodergren E, Weinstock GM, Botstein D, Desai MM (2013) Pervasive genetic hitchhiking and clonal interference in forty evolving yeast populations. *Nature* 500(7464):571–574. <https://doi.org/10.1038/nature12344>
- Lauring AS (2020) Within-host viral diversity: a window into viral evolution. *Ann Rev Virol* 7:63–81. <https://doi.org/10.1146/annurev-virology-010320-061642>. <https://www.annualreviews.org/content/journals/10.1146/annurev-virology-010320-061642>
- Lipman DJ, Wilbur WJ (1991) Modelling neutral and selective evolution of protein folding. *Proc Roy Soc London B* 245:7–11
- Louis AA (2016) Contingency, convergence and hyper-astronomical numbers in biological evolution. *Stud Hist Philos Sci C* 58:107–116
- Manrubia S (2022) The simple emergence of complex molecular function. *Phil Trans R Soc A Math Phys Eng Sci* 380(2227):20200422. <https://doi.org/10.1098/rsta.2020.0422>
- Manrubia S, Cuesta JA (2015) Evolution on genotype networks accelerates the ticking rate of the molecular clock. *J R Soc Interface* 12:20141010
- Manrubia S, Cuesta JA (2017) Distribution of genotype network sizes in sequence-to-structure genotype-phenotype maps. *J R Soc Interface* 14:20160976
- Manrubia S, Cuesta JA, Aguirre J, Ahnert SE, Altenberg L, Cano AV, Catalán P, Diaz-Uriarte R, Elena SF, Garcia-Martin JA, Hogeweg P, Khatri BS, Krug J, Louis AA, Martin NS, Payne JL, Tarnowski MJ, Weiß M (2021) From genotypes to organisms: state-of-the-art and perspectives of a cornerstone in evolutionary dynamics. *Phys Life Rev* 38:55–106. <https://doi.org/10.1016/j.pprev.2021.03.004>
- Martin NS, Ahnert SE (2022) Fast free-energy-based neutral set size estimates for the rna genotype-phenotype map. *J R Soc Interface* 19(191):20220072. <https://doi.org/10.1098/rsif.2022.0072>
- Martin NS, Camargo CQ, Louis AA (2024) Bias in the arrival of variation can dominate over natural selection in richard dawkins's biomorphs. *PLOS Comput Biol* 20(3):1–31. <https://doi.org/10.1371/journal.pcbi.1011893>
- Maynard Smith J (1970) Natural selection and the concept of a protein space. *Nature* 225:563–564
- Nowak S, Krug J (2013) Accessibility percolation on the hypercube and the phase transition of greedy adaptive walks. *Europhys Lett* 101:66004. <https://doi.org/10.1209/0295-5075/101/66004>
- Ohta T (1973) Slightly deleterious mutant substitutions in evolution. *Nature* 246:96–98
- Piacentini L, Fanti L, Specchia V, Bozzetti MP, Berloco M, Palumbo G, Pimpinelli S (2014) Transposons, environmental changes, and heritable induced phenotypic variability. *Chromosoma* 123(4):345–354. <https://doi.org/10.1007/s00412-014-0464-y>
- Pigliucci M (2012) Landscapes, surfaces, and morphospaces: what are they good for? In: Svensson E, Calsbeek R (eds) *The adaptive landscape in evolutionary biology*, p 26. Oxford University Press, Oxford
- Schaper S, Louis AA (2014) The arrival of the frequent: how bias in genotype-phenotype maps can steer populations to local optima. *PLoS ONE* 9:e86635

- Schmiegelt B, Krug J (2023) Accessibility percolation on cartesian power graphs. *J Math Biol* 86: 46. <https://doi.org/10.1007/s00285-023-01882-z>
- Schultes EA, Bartel DP (2000) One sequence, two ribozymes: implications for the emergence of new ribozyme folds. *Science* 289:448–452
- Schuster P, Sigmund K (1983) Replicator dynamics. *J Theor Biol* 100:533–538
- Schuster P, Swetina J (1988) Stationary mutant distributions and evolutionary optimization. *Bull Math Biol* 50:635–660
- Schuster P, Fontana W, Stadler PF, Hofacker IL (1994) From sequences to shapes and back: a case study in RNA secondary structures. *Proc Roy Soc London B* 255:279–284
- Seoane LF, Secaira-Morocho H, Somovilla P, Lázaro E, Manrubia S (2026) Hierarchical genotype networks and incipient speciation in Q β phage quasispecies. *Proc Natl Acad Sci* 123(14) e2512150123. <https://doi.org/10.1073/pnas.2512150123>
- Sinzelle L, Izsvák Z, Ivics Z (2009) Molecular domestication of transposable elements: from detrimental parasites to useful host genes. *Cellular Mol Life Sci* 66:1073–93. <https://doi.org/10.1007/s00018-009-8376-3>
- Sun H, Liu H, Zhang J, Qu X, Pang Z, Xu F, Wu C, Jiang Y, Shi M, Liu Q, Liao M (2025) Genome-scale evolution and phylodynamics of swine influenza A viruses in China: a genomic epidemiology study. *Lancet Microbe* 6(6):101020. <https://doi.org/10.1016/j.lanmic.2024.101020>
- Symons RH (1981) Avocado sunblotch viroid: primary sequence and proposed secondary structure. *Nucleic Acids Res* 9(23):6527–6537. <https://doi.org/10.1093/nar/9.23.6527>
- Szendro IG, Schenk MF, Franke J, Krug J, de Visser JAGM (2013) Quantitative analyses of empirical fitness landscapes. *J Stat Mech Theory Exp* 2013(01):P01005. <https://doi.org/10.1088/1742-5468/2013/01/P01005>
- Vakirlis N, Carvunis AR, McLysaght A (2020) Synteny-based analyses indicate that sequence divergence is not the main source of orphan genes. *eLife* 9:e53500. <https://doi.org/10.7554/elife.53500>
- Vignuzzi M, Stone JK, Arnold JJ, Cameron CE, Andino R (2006) Quasispecies diversity determines pathogenesis through cooperative interactions in a viral population. *Nature* 439(7074): 344–348. <https://doi.org/10.1038/nature04388>
- Villanueva A, Secaira-Morocho H, Seoane LF, Lázaro E, Manrubia S (2022) Genotype-to-protein map and collective adaptation in a viral population. *Biophysica* 2(4):381–399. <https://doi.org/10.3390/biophysica2040034>. <https://www.mdpi.com/2673-4125/2/4/34>
- Viver T, Conrad RE, Rodríguez-R LM, Ramírez AS, Venter SN, Rocha-Cárdenas J, Llabrés M, Amann R, Konstantinidis KT, Rossello-Mora R (2024) Towards estimating the number of strains that make up a natural bacterial population. *Nat Commun* 15(1):544. <https://doi.org/10.1038/s41467-023-44622-z>
- Wachowius F, Porebski BT, Johnson CM, Holliger P (2021) Emergence of function from single RNA sequences by darwinian evolution. *bioRxiv*. <https://doi.org/10.1101/2021.03.03.433769>
- Wilke CO, Wang JL, Ofria C, Lenski RE, Adami C (2001) Evolution of digital organisms at high mutation rates leads to survival of the flattest. *Nature* 412:331–333
- Wolff A, Krug J (2009) Robustness and epistasis in mutation-selection models. *Phys Biol* 6(3): 036007
- Wright S (1931) Evolution in Mendelian populations. *Genetics* 16:97–159
- Yona AH, Alm EJ, Gore J (2018) Random sequences rapidly evolve into de novo promoters. *Nat Commun* 9(1):1530
- Zagorski M, Burda Z, Waclaw B (2016) Beyond the hypercube: evolutionary accessibility of fitness landscapes with realistic mutational networks. *PLoS Comput Biol* 12:e1005218
- Zanini F, Puller V, Brodin J, Albert J, Neher RA (2017) In vivo mutation rates and the landscape of fitness costs of HIV-1. *Virus Evol* 3(1):vex003