

Mount Fuji's stubby peak: the genotypic density of additive landscapes near maximal fitness

Justin B. Kinney  *

Simons Center for Quantitative Biology, Cold Spring Harbor Laboratory, Cold Spring Harbor, New York 11724, United States

*Corresponding author: Simons Center for Quantitative Biology, Cold Spring Harbor Laboratory, 1 Bungtown Rd., Cold Spring Harbor, New York 11724, United States.
Email: jkinney@cshl.edu.

Additive fitness landscapes—also called Mount Fuji landscapes—are the simplest and most widely used models of sequence-function relationships. As such, they play essential roles across multiple areas of biology, including evolutionary theory, quantitative genetics, gene regulation, and protein science. One of the most basic properties of any fitness landscape is its *genotypic density*—the number of sequences near a given fitness value. Understanding this density is especially important near fitness peaks, as it quantifies the supply of high-fitness genotypes. Here I study the genotypic density of additive landscapes near fitness peaks. Although this density is well known to be approximately Gaussian near the middle of the fitness range, its behavior near maximal fitness has not been reported. I begin by deriving a saddle-point approximation that accurately describes the genotypic density of additive landscapes over virtually the entire fitness range. I then show that the log of this density follows a power law near maximal fitness, with the exponent determined by how much the best allele at each position outperforms its nearest competitor. This power-law behavior holds over a substantial fraction of fitness values, besting the Gaussian approximation on both simulated and empirical landscapes across roughly a quarter to a third of the fitness range. Under certain conditions this behavior also extends to globally epistatic landscapes (defined as nonlinear functions of one or more additive traits), though with a reduced range of validity. These findings advance our understanding of one of the most fundamental models of sequence-function relationships. In particular, they reveal that the uppermost reaches of Mount Fuji landscapes, rather than being sharply peaked, are actually quite stubby.

Keywords: fitness landscapes; additive models; genotypic density; saddle-point approximation; global epistasis

Introduction

Fitness landscapes—quantitative mappings from biological sequences to fitness values—have long been a key organizing concept in evolutionary biology (Wright 1932; Maynard Smith 1970; Fragata et al. 2019). The simplest and most widely used model for such landscapes is the *additive model*, in which each position in a sequence contributes independently to fitness. Additive models are used in many areas of biology, both in theoretical efforts and in empirical studies of sequence-function relationships. In particular, they form the basis for *global epistasis models*, which are defined as nonlinear functions of one or more additive traits. Many experimentally measured fitness landscapes are well-described by either additive or globally epistatic models (Carneiro and Hartl 2010; Kinney et al. 2010; Kryazhimskiy et al. 2014; Stiffler et al. 2015; Adams et al. 2016; Sarkisyan et al. 2016; Sailer and Harms 2017; Otwinowski 2018; Otwinowski et al. 2018; Baeza-Centurion et al. 2019; Starr et al. 2020; Faure et al. 2024). A thorough understanding of the quantitative properties of additive models is therefore warranted.

One of the simplest questions one can ask about a fitness landscape is how many sequences have fitness near a given value. I refer to this quantity as *genotypic density*. In the case of additive landscapes, the Central Limit Theorem (CLT) suggests that the genotypic density is approximately Gaussian in the bulk of the

fitness distribution. However, this bulk approximation is qualitatively wrong near the maximal achievable fitness. Although the importance of genotypic density has been recognized by many authors, both for modeling evolution (Berg and von Hippel 1987; Aita and Husimi 1996; Sengupta et al. 2002; Djordjevic et al. 2003; Berg et al. 2004; Mustonen and Lassig 2005; Sella and Hirsh 2005; Lässig 2007; Mustonen et al. 2008; Neher and Shraiman 2011) and for bioinformatic analyses (Staden 1989; Touzet and Varré 2007; Madsen et al. 2017), the quantitative form of the near-peak density of additive fitness landscapes has not been reported.

In the fitness landscape literature, additive landscapes are sometimes known as *Mount Fuji landscapes*, a moniker chosen to emphasize their simple topographic structure: a single peak that is evolutionarily accessible from any starting sequence (Aita and Husimi 1996, 2000; Aita et al. 2004). Although the Gaussian behavior of genotypic density in the bulk was realized early on in this context (Aita and Husimi 1996), the behavior near maximal fitness appears not to have been investigated. This is perhaps because additive landscapes have primarily been considered as null models against which to assess more complex epistatic models (including NK, Rough Mount Fuji, and House of Cards models) (Szendro et al. 2013; Neidhart et al. 2014; Pahujani and Krug 2025), rather than as objects of primary scientific interest. The near-peak genotypic density nevertheless has evolutionary implications: it quantifies the supply of high-fitness genotypes

accessible by mutation, and therefore shapes both adaptive and purifying selection.

In quantitative genetics, additive models have long provided a theoretical foundation for the study of plant and animal breeding (Barton and Keightley 2002; Walsh and Lynch 2018). Of particular importance is the infinitesimal model, which assumes that a trait is determined by the sum of effects across a large number of loci (Barton et al. 2017). Both theory and data indicate that additive genetic contributions account for most of the heritable variation in complex traits, even when epistasis is present at individual loci (Hill et al. 2008). The distribution of trait values in a population under selection has been studied extensively, including using concepts and methods from statistical physics (Barton and Coe 2009; Barton and de Vladar 2009). In particular, the departure of population trait distributions from Gaussianity has been explored (Turelli and Barton 1994). Genotypic density itself, however, does not appear to have received a similar treatment.

In the bioinformatic study of gene regulation, additive landscapes are commonly used to model the binding of transcription factors (TFs) to DNA. Additive models such as position weight matrices (PWMs) are routinely used to scan genomes for putative TF binding sites (Stormo and Fields 1998; Stormo 2000; Djordjevic et al. 2003; Foat et al. 2006; Elemento et al. 2007; Kinney et al. 2007; Grant et al. 2011), and large databases of such models are widely used (Wingender et al. 2000; Pachkov et al. 2007; Vorontsov et al. 2024; Baydar et al. 2025). Assessing the statistical significance of high-scoring sites requires accurate estimates of the upper tail of the genotypic density, and a few studies have developed methods for doing this (Staden 1989; Touzet and Varré 2007; Madsen et al. 2017). Of particular note is Madsen et al. (2017), who developed an estimate of p -values for PWMs using a saddle-point approximation similar to the one I use below. However, neither Madsen et al. (2017) nor other work in this area has described the mathematical form of the near-peak density.

Beginning with foundational studies by Berg and von Hippel (Berg and von Hippel 1987, 1988), researchers at the interface of statistical physics and population genetics have used additive models for TF binding energy to study the evolution of TF binding sites (Gerland et al. 2002; Sengupta et al. 2002; Djordjevic et al. 2003; Sella and Hirsh 2005; Neher and Shraiman 2011). Genotypic density plays an important role in this literature. When evolution is modeled in the weak-mutation limit, the equilibrium distribution of binding site sequences takes the form of a canonical ensemble from statistical physics, one in which fitness values are weighted by an exponentially tilted form of the genotypic density (Berg et al. 2004; Mustonen and Lassig 2005; Lässig 2007; Mustonen et al. 2008). The original work of Berg and von Hippel derives the exact genotypic density for an idealized additive model in which every mutation away from the highest-fitness sequence produces the same fitness deficit. To treat more general additive models, Sengupta et al. (2002) derived a saddle-point estimate of genotypic density (the same estimate I consider below), but used it only to recover the Gaussian approximation suggested by the CLT.

Here I study the genotypic density of additive landscapes near maximal fitness. I first derive a saddle-point approximation to this density and show that it is remarkably accurate across nearly the entire fitness range. I then investigate the near-peak behavior of this approximation. I find that, for an important class of idealized models, the log density follows a power law (or equivalently, the density itself has an exponentiated power-law form). Additional considerations as well as empirical results suggest that many real-world landscapes also follow an exponentiated power law.

A consequence is that, as fitness decreases from its maximum, the number of available genotypes initially expands very rapidly, then grows more and more slowly until finally it matches up with the predictions of the Gaussian approximation. The landscape is therefore not sharply peaked at its maximum but instead is broad and gently rounded—more like the top of a rolling hill than the summit of a mountain. This motivates my use of the term “stubby.”

I validate these findings on simulated landscapes and on two empirical landscapes: one from a massively parallel reporter assay on a transcriptional promoter (Kinney et al. 2010), and one from a deep mutational scanning experiment on a protein domain (Olson et al. 2014). In all cases these landscapes exhibit stubby peaks that are well described by an exponentiated power law. I then show that globally epistatic landscapes defined on one or more traits often exhibit near-peak densities of the same mathematical form, though with caveats. The derivations of these results are provided below. Additional mathematical details are given in Supplemental Information (SI).

Results

Preliminaries

Consider sequences of length L comprising characters from an alphabet of size C (e.g. $C = 4$ for DNA or RNA, $C = 20$ for proteins). Each sequence x is represented by a one-hot encoded vector $x = \{x_{lc}\}$, where $x_{lc} = 1$ if character c occurs at position l and $x_{lc} = 0$ otherwise. Additive fitness landscapes have the form

$$f(x) = \theta_0 + \sum_{l,c} \theta_{lc} x_{lc}, \quad (1)$$

where θ_{lc} quantifies the effect of character c at position l and θ_0 is an overall fitness baseline. In what follows I denote the maximal fitness effect at position l by $\theta_l^{\max} = \max_c \theta_{lc}$, the optimal character by $c_l^{\max}$, and the fitness-maximizing sequence by $x_{\max}$. The maximum achievable fitness is $F_{\max} = f(x_{\max})$; $F_{\min}$ denotes the minimum achievable fitness.

Throughout this paper I illustrate the claims and findings using additive models for two empirical landscapes (Fig. 1). The first model describes the transcriptional activity of variants of the *Escherichia coli lac* promoter ($L = 75$, $C = 4$) measured using a massively parallel reporter assay (MPRA) based on fluorescence-activated cell sorting (Kinney et al. 2010). The second model describes the binding of variants of the B1 domain of protein G (GB1; $L = 55$, $C = 20$) to IgG, as measured using a deep mutational scanning (DMS) experiment based on mRNA display (Olson et al. 2014). Additive models for both landscapes were inferred from experimental data using MAVE-NN (Tareen et al. 2022), which separately accounts for any global epistasis nonlinearities present in the data.

Defining the genotypic density

The focus of this paper is the quantitative form of the genotypic density, $\rho(F)$, which is defined so that $\rho(F) dF$ is the number of sequences within an infinitesimal interval $[F, F + dF]$. This density is given in terms of Dirac delta functions by

$$\rho(F) = \sum_x \delta(F - f(x)). \quad (2)$$

Below I develop various smooth mathematical approximations to $\rho(F)$. The reader might reasonably question what it means to approximate $\rho(F)$ with a smooth function, since no nonzero smooth




Fig. 1. Two empirical landscapes analyzed in this paper. (a) Additive model for the *lac* promoter landscape of Kinney et al. (2010) ($L = 75$, $C = 4$), visualized as a sequence logo (Tareen and Kinney 2019). Positions are relative to the transcription start site. (b) Additive model for the GB1 landscape of Olson et al. (2014) ($L = 55$, $C = 20$), visualized as a heatmap. Positions correspond to residues 2–56. Dots mark wild-type residues. The parameters of both models were inferred from data using MAVE-NN (Tareen et al. 2022), then shifted and scaled so that the fitness values represent z-scores, i.e. random sequences have fitness values with zero mean and unit variance.

function will provide a good approximation in the standard squared-error sense.

Perhaps the most intuitive reason for considering smoothed approximations to $\rho(F)$ is that, in reality, fitness values are never measured with infinite precision. One way to account for this is to replace the delta function in Equation (2) with a Gaussian distribution that has a standard deviation commensurate with the expected level of experimentally unresolvable uncertainty. Doing this yields a kernel density estimate that is smooth by construction. In practice, the form of this estimate is almost completely independent of the width of the Gaussian as long as the uncertainty scale is much less than the fitness range $F_{\max} - F_{\min}$, and much greater than the gaps between the extremal fitness values ($F_{\max}$ and $F_{\min}$) and their nearest neighbors. This condition is often satisfied for a wide range of uncertainties in practice.

The bulk approximation

In an additive landscape, the fitness of each sequence is a sum of L independent contributions. The CLT therefore suggests that F values are approximately normally distributed, and thus that $\rho(F)$ can be approximated by

$$\rho_{\text{bulk}}(F) = \frac{C^L}{\sqrt{2\pi\sigma_{\text{bulk}}^2}} \exp\left[-\frac{(F - F_{\text{mean}})^2}{2\sigma_{\text{bulk}}^2}\right], \quad (3)$$

where F_{mean} and σ_{bulk}^2 are the mean and variance of the fitness values observed for randomly generated sequences. The CLT, however, only guarantees accuracy within a few standard deviations of the mean. Indeed, Fig. 2 shows that this Gaussian form accurately describes the central mass of $\rho(F)$, but breaks down dramatically near $F_{\max}$. Specifically, $\rho_{\text{bulk}}(F)$ provides increasingly inflated estimates of $\rho(F)$ as F approaches $F_{\max}$ from below, and predicts nonzero density at all fitness values greater than $F_{\max}$, where the true density vanishes. The breakdown worsens as sequences grow longer. Assuming the distribution of θ_{lc} values remains fixed, the width of the fitness distribution scales as $\sigma_{\text{bulk}} \propto \sqrt{L}$, whereas $F_{\max} - F_{\text{mean}}$ grows linearly with L . The fitness peak therefore recedes to ever more extreme quantiles as L increases, placing it in a regime where the CLT provides no useful information. To characterize $\rho(F)$ across the full fitness range—and in particular near $F_{\max}$ —a fundamentally different approach is required.

The tilted distribution and free fitness

To derive an improved estimate of $\rho(F)$, we first introduce a family of exponentially tilted distributions over sequence space:

$$p_{\beta}(x) = e^{\beta f(x) - \Phi(\beta)}, \quad (4)$$

where β is a scalar that parameterizes the family, and

$$\Phi(\beta) = \log \sum_x e^{\beta f(x)} \quad (5)$$

is a function that serves to normalize the distribution. By varying β from $-\infty$ to ∞ we can center this distribution over specific fitness values. In particular, $\beta = 0$ recovers the uniform distribution on sequence space (which has fitness values concentrated near F_{mean}), while sufficiently large β concentrates the distribution near $F_{\max}$.

Distributions of the form in Equation (4) arise naturally in both population genetics and statistical physics. In the weak-mutation limit of Wright-Fisher processes with reversible mutation rates, $p_{\beta}(x)$ describes the stationary distribution over genotypes, where β is related to the effective population size via $\beta = 2N_e$ (Berg et al. 2004; Mustonen and Lassig 2005; Sella and Hirsh 2005; Lässig 2007; Neher and Shraiman 2011). If instead we interpret x as the microstate of a physical system and $-f(x)$ as the corresponding energy, then $p_{\beta}(x)$ describes the distribution over microstates that arises when the system is at temperature $1/\beta$. In such contexts p_{β} is called the *canonical ensemble*.

The normalizing term $\Phi(\beta)$ encodes important information about the tilted distribution. Indeed, this quantity has previously been termed *free fitness* (Iwasa 1988; Sella and Hirsh 2005; Lässig 2007; Barton and de Vladar 2009), as it is analogous to negative the free energy in physical systems. One key property is that $\Phi(\beta)$ is a cumulant-generating function for the marginal distribution of fitness values,

$$p_{\beta}(F) = \sum_x p_{\beta}(x) \delta(F - f(x)). \quad (6)$$

In other words, the mean of $p_{\beta}(F)$ is $\mu_{\beta} = \Phi'(\beta)$, the variance is $\sigma_{\beta}^2 = \Phi''(\beta)$, and so on. This is true regardless of the functional form of $f(x)$.




Fig. 2. Accuracy of the saddle-point approximation on simulated and empirical landscapes. (a,b) Results for simulated landscapes having Gaussian effects θ_i and alphabets of different size. (c) Results for the *lac* promoter landscape of Kinney et al. (2010). (d) Results for the GB1 landscape of Olson et al. (2014). Fitness values for the empirical landscapes correspond to z-scores, as in Fig. 1. Each panel compares the saddle-point approximation ρ_{saddle} (orange solid line) to both the Gaussian approximation ρ_{bulk} (olive dashed line) and the true density ρ_{true} (black solid line), computed using the algorithm of Touzet and Varré (2007). Vertical dotted lines indicate F_{max} and F_{mean} .

In the specific case where $f(x)$ is additive (Equation (1)), free fitness can be written in closed form:

$$\Phi(\beta) = \beta\theta_0 + \sum_i \log \sum_c e^{\beta\theta_{ic}}. \quad (7)$$

Corresponding expressions for μ_β and σ_β^2 are given in SI Sec. S1. It is readily shown that the higher-order cumulants become irrelevant as L grows, and thus that $p_\beta(F)$ becomes well-approximated by the Gaussian distribution

$$p_\beta(F) = \frac{1}{\sqrt{2\pi\sigma_\beta^2}} \exp\left[-\frac{(F - \mu_\beta)^2}{2\sigma_\beta^2}\right]. \quad (8)$$

The saddle-point approximation

We now return to the problem of estimating the genotypic density $\rho(F)$ in Equation (2). Since the weight for each sequence in the tilted distribution depends on x only through $f(x)$, it can be taken outside the sum in Equation (6). Additionally using Equation (2) and Equation (4) then gives

$$p_\beta(F) = e^{\beta F - \Phi(\beta)} \rho(F). \quad (9)$$

We now have two distinct expressions for $p_\beta(F)$, Equations (8) and (9). Equating these and solving for $\rho(F)$ yields a candidate approximation for genotypic density

$$\rho_\beta(F) = \frac{e^{\Phi(\beta) - \beta F}}{\sqrt{2\pi\sigma_\beta^2}} \exp\left[-\frac{(F - \mu_\beta)^2}{2\sigma_\beta^2}\right]. \quad (10)$$

We are still free to choose β however we like. We expect, however, that $\rho_\beta(F)$ will be a good approximation only when F is within a few standard deviations σ_β of the mean μ_β . Indeed, Equation (10) reduces to the bulk approximation Equation (3) if we set $\beta = 0$, and the entire motivation for this paper is that $\rho_{\text{bulk}}(F)$ is a poor estimate of the true density $\rho(F)$ when F is near F_{max} .

The CLT suggests that the approximation in Equation (8) is most accurate when $\mu_\beta = F$. We therefore choose β to be a function of F defined implicitly by this requirement, i.e.

$$\rho_{\text{saddle}}(F) = \frac{e^{\Phi(\beta) - \beta F}}{\sqrt{2\pi\sigma_\beta^2}} \quad \text{where } \beta \text{ is s.t. } \Phi'(\beta) = F. \quad (11)$$

This is known as the *saddle-point* approximation (Darwin and Fowler 1922; Daniels 1954; Reid 1988; Barndorff-Nielsen and Cox 1989). The name comes from an alternative derivation in which one approximates the value of a contour integral in the complex plane by its value near a stationary point of the integrand (i.e. a saddle-point). The above derivation, however, more clearly explains why the saddle-point approximation works so well throughout the entire fitness range. See Reid (1988) for further discussion of these two derivations. Equation (11) can also be thought

of in terms of large deviation theory (Touchette 2009), with rate function given by $[\beta F - \Phi(\beta)]/L$. The above derivation, however, has the advantage of also providing the $1/\sqrt{2\pi\sigma_\beta^2}$ prefactor, which is generally neglected in rate function calculations.

Performance of the saddle-point approximation

To assess the performance of ρ_{saddle} , I compared its predictions to the true densities of two simulated landscapes and of the two empirical landscapes mentioned earlier. The results are shown in Fig. 2. In all cases ρ_{saddle} closely tracks ρ over the entire fitness range. By contrast, ρ_{bulk} progressively overestimates ρ as F approaches F_{max} , ultimately deviating by many orders of magnitude. This agreement holds equally for simulated and empirical landscapes over both DNA/RNA ($C = 4$) and protein ($C = 20$) alphabets.

I note that computing the true density ρ is complicated by the fact that all four landscapes in Fig. 2 are too large to exhaustively compute the fitness values of all sequences. I therefore used the dynamic programming algorithm of Touzet and Varré (2007) instead. This algorithm computes cumulative counts above a given threshold by iteratively refining upper and lower bounds using discretized fitness values. Running this algorithm to full precision across the entire fitness range is prohibitive. Restricting it to the upper 30% of the range and terminating refinement before full convergence, however, yields certified upper and lower bounds that are visually indistinguishable in Fig. 2.

Near-peak scaling for binary alphabets

Equation (11) does not by itself reveal the mathematical form of genotypic density near maximal fitness. What we require is an explicit expression for $\rho(F)$ near F_{max} , one that is not obscured by the implicit dependence of β on F . To this end, we express the free fitness as $\Phi(\beta) = \beta F_{\text{max}} + LE(\beta)$, where

$$E(\beta) = \frac{1}{L} \sum_{l=1}^L \log \left(1 + \sum_{c \neq c_l^{\text{max}}} e^{-\beta \delta_{lc}} \right), \quad \delta_{lc} = \theta_l^{\text{max}} - \theta_{lc}. \quad (12)$$

We call this the *per-site deficit contribution* to free fitness, and define the *per-site fitness deficit* as $\epsilon = (F_{\text{max}} - F)/L$. Substituting these into Equation (11) and dropping the prefactor (which is sub-leading for large L) gives

$$\frac{1}{L} \log \rho_{\text{saddle}} \approx \beta \epsilon + E(\beta), \quad (13)$$

where β is an implicit function of ϵ determined by $E'(\beta) = -\epsilon$. Since every $\delta_{lc} > 0$ in Equation (12), $E(\beta) \rightarrow 0$ as $\beta \rightarrow \infty$. However, the near-peak behavior of ρ_{saddle} depends on precisely how $E(\beta)$ approaches this limit. We therefore seek to determine the leading-order behavior of $E(\beta)$ as $\beta \rightarrow \infty$.

To simplify this analysis, we first consider binary alphabets ($C = 2$). In this case, the deficit contribution becomes

$$E_{\text{bin}}(\beta) = \frac{1}{L} \sum_{l=1}^L \log(1 + e^{-\beta \Delta_l}), \quad (14)$$

where Δ_l is the (positive) fitness gap between the two additive effects at position l . Replacing the empirical distribution of these gaps with a smooth density $p_{\text{gap}}(\Delta)$ gives

$$E_{\text{bin}}(\beta) \approx \int \log(1 + e^{-\beta \Delta}) p_{\text{gap}}(\Delta) d\Delta = \frac{1}{\beta} \int K(s) p_{\text{gap}}\left(\frac{s}{\beta}\right) ds, \quad (15)$$

where we have substituted $s = \beta \Delta$ and defined $K(s) = \log(1 + e^{-s})$. It is helpful to think of $K(s)$ as a smoothing kernel: it equals $\log 2$ at the origin, has width of order 1, decays as e^{-s} for $s \gg 1$, and has area $\frac{\pi^2}{12}$ (see SI Sec. S2.1). Since larger β probes $p_{\text{gap}}(\frac{s}{\beta})$ at ever-smaller arguments, only the behavior of $p_{\text{gap}}(\Delta)$ near $\Delta = 0$ matters.

When p_{gap} is *regular* (by this I mean $0 < p_{\text{gap}}(0) < \infty$), Equation (15) can be evaluated to leading order:

$$E_{\text{bin}}(\beta) \approx \frac{\pi^2 p_{\text{gap}}(0)}{12\beta} + O(\beta^{-2}). \quad (16)$$

The saddle-point condition then gives $\epsilon = \frac{\pi^2 p_{\text{gap}}(0)}{12\beta^2}$, which inverts to $\beta = \sqrt{\frac{\pi^2 p_{\text{gap}}(0)}{12\epsilon}}$. Substituting Equation (16) and this expression for β back into Equation (13) yields our desired result:

$$\frac{1}{L} \log \rho \approx B \epsilon^{1/2}, \quad \text{where } B = \sqrt{\frac{\pi^2 p_{\text{gap}}(0)}{3}}. \quad (17)$$

It is tempting to assume that p_{gap} is always regular, and indeed it is if each θ_{lc} parameter is drawn i.i.d. from a sufficiently well-behaved distribution (see SI Sec. S3). Regularity is not guaranteed, however, when the magnitude of additive effects varies across positions. For example, if most positions in a sequence have little to no effect, the gap distribution may scale as Δ^γ near $\Delta = 0$ for some negative exponent γ . (Note that $\gamma > -1$ is required for p_{gap} to be integrable near $\Delta = 0$). This is an important case to consider because many empirical landscapes exhibit this type of behavior.

We therefore consider the more general case where $p_{\text{gap}}(\Delta) \sim c \Delta^\gamma$ near $\Delta = 0$, for some exponent $\gamma > -1$ and coefficient $c > 0$. Repeating the asymptotic analysis with this assumption yields a more general power law:

$$\frac{1}{L} \log \rho \approx B \epsilon^\alpha, \quad \text{where } \alpha = \frac{1+\gamma}{2+\gamma}, \quad (18)$$

and B is a constant that depends on γ and c ; see SI Sec. S2 for a full derivation. Note that the exponent α increases monotonically with γ . For regular gap distributions ($\gamma = 0$), one recovers $\sqrt{\epsilon}$ scaling ($\alpha = \frac{1}{2}$). When the gap distribution is enriched for small gaps ($\gamma < 0$), as is expected when most positions have small fitness effects, the density falls off more quickly near F_{max} than $\sqrt{\epsilon}$ scaling predicts, producing a stubbier peak ($0 < \alpha < \frac{1}{2}$). Conversely, a deficit of small gaps ($\gamma > 0$) causes the density to fall off more slowly, producing a sharper peak ($\frac{1}{2} < \alpha < 1$).

In the extreme case where all gaps are bounded below by some $\Delta_0 > 0$, $\log \rho$ grows approximately linearly with ϵ (i.e. $\alpha = 1$), but is also tempered by a log-linear term in ϵ that reduces this growth away from $\epsilon = 0$. This is exemplified by the landscape proposed by Berg and von Hippel (1987), in which $F_{\text{max}} - f(x)$ is Δ_0 times the Hamming distance between x and x_{max} , and consequently $p_{\text{gap}}(\Delta) = \delta(\Delta - \Delta_0)$. See SI Sec. S4 for details.

Figure 3 validates these predictions for binary landscapes with four choices of gap distribution: regular ($\gamma = 0$), enriched for small gaps ($\gamma < 0$), depleted of small gaps ($\gamma > 0$), and fixed gap size (Hamming-distance landscape). In all cases, the near-peak approximation closely tracks ρ_{saddle} near F_{max} , while ρ_{bulk} progressively overestimates the density. The same curves plotted against ϵ^α are shown in SI Supplementary Fig. S1.




Fig. 3. Near-peak scaling for additive landscapes on binary sequences ($C = 2, L = 200$). Each panel compares the saddle-point approximation ρ_{saddle} (orange solid line) to the Gaussian approximation ρ_{bulk} (olive dashed line) and the near-peak approximation ρ_{peak} (blue dashed line). (a) Gaussian fitness effects ($\theta_{lc} \sim \mathcal{N}(0, 1)$), producing a regular gap distribution with $\gamma = 0$ and $\alpha = 1/2$. (b) Left-skewed gap density ($p_{\text{gap}} \propto \Delta^{-1/2}$, $\gamma = -1/2$), enriched for small gaps and producing a stubbier peak ($\alpha = 1/3$). (c) Right-skewed gap density ($p_{\text{gap}} \propto \Delta$, $\gamma = 1$), depleted of small gaps and producing a sharper peak ($\alpha = 2/3$). (d) Hamming-distance landscape with uniform gap $\Delta = 1$ at all positions ($\alpha = 1$ with log-linear correction). Vertical dotted lines indicate F_{max} and F_{mean} . See [Supplementary Fig. S1](#) for plots made with respect to ϵ^α instead of F .

Near-peak scaling for general alphabets

The preceding analysis was restricted to binary alphabets, where each position contributes a single fitness gap Δ_l to $E(\beta)$. This allowed $E(\beta)$ to be evaluated in the large- L limit by replacing the sum over positions with an integral over the gap density $p_{\text{gap}}(\Delta)$. For alphabets with $C \geq 3$ characters, however, $E(\beta)$ in Equation (12) includes contributions from multiple nonoptimal characters at each position. The resulting multidimensional integral cannot be evaluated so easily.

Fortunately, we can bound $E(\beta)$ both above and below using the binary alphabet results. Define Δ_l to be the gap between the additive effects of the optimal and runner-up characters at position l , and let $E_{\text{bin}}(\beta)$ be the deficit contribution obtained by restricting the alphabet to only these two characters. For the lower bound, this restriction removes terms from the sum inside the logarithm in Equation (12), which can only decrease $E(\beta)$, giving $E_{\text{bin}}(\beta) \leq E(\beta)$. For the upper bound, we use $\Delta_l \leq \delta_{lc}$ for all $c \neq c_l^{\text{max}}$, together with the subadditivity of $\log(1+x)$, to get

$$E(\beta) \leq \frac{1}{L} \sum_l \log[1 + (C-1)e^{-\beta\Delta_l}] \leq \frac{C-1}{L} \sum_l \log[1 + e^{-\beta\Delta_l}]. \quad (19)$$

We thus obtain the sandwich bound,

$$E_{\text{bin}}(\beta) \leq E(\beta) \leq (C-1)E_{\text{bin}}(\beta). \quad (20)$$

This implies that $E(\beta)$ exhibits the same asymptotic behavior in the $\beta \rightarrow \infty$ limit as its binary alphabet counterpart. In particular, when the runner-up gap distribution satisfies $p_{\text{gap}}(\Delta) \sim c\Delta^\gamma$ near the origin, the scaling exponent $\alpha = (1+\gamma)/(2+\gamma)$ carries over unchanged from the binary case, though the constant B may differ.

Accurately describing real-world landscapes requires one additional consideration. If all C characters at a position make the same contribution to fitness ($\delta_{lc} = 0$ for all c), that position contributes a constant $\frac{1}{L} \log C$ to $E(\beta)$ that is independent of β . The cumulative effect of such positions is to add a constant to $\frac{1}{L} \log \rho$. Similar but smaller contributions arise at positions where some but not all characters are tied for maximal fitness. These considerations motivate the following hypothesis: that the near-peak genotypic density of additive landscapes will often follow an exponentiated power law of the form

$$\frac{1}{L} \log \rho \approx A + B\epsilon^\alpha, \quad (21)$$

where the specific values of A , B , and α for each landscape can be determined by fitting this function to the saddle-point approximation in Equation (11).

I propose Equation (21) as an effective description of many additive landscapes. I emphasize, though, that this form is not exact and will not accurately describe every landscape. Rather, it is a parametric approximation that is motivated by idealized models




Fig. 4. Near-peak scaling for general alphabets, tested on the same four landscapes as in Fig. 2. Each panel plots $\log \rho(F)$ versus F , comparing the saddle-point approximation ρ_{saddle} , the Gaussian approximation ρ_{bulk} , and the near-peak scaling form ρ_{peak} , the parameters of which (A , B , and α) were obtained by fitting $L^{-1} \log \rho = A + B \epsilon^a$ to the saddle-point approximation. Empirical F values are measured in the z-score units defined in Fig. 1. The vertical dotted line at F_{cross} marks the crossover fitness at which ρ_{peak} becomes a closer approximation to ρ_{saddle} than ρ_{bulk} , with the crossover ratio r labeled. (a) Simulated landscape using $C = 4$. (b) Simulated landscape using $C = 20$. (c) *lac* promoter landscape ($L = 75$, $C = 4$). (d) GB1 protein landscape ($L = 55$, $C = 20$). See Supplementary Fig. S2 for plots made with respect to ϵ^a instead of F .

and that appears to often work well in practice (including for landscapes with gap distributions that do not follow a power law).

Figure 4 tests this hypothesis on the same four landscapes considered in Fig. 2: two simulated (i.i.d. Gaussian effects with $L = 100$, $C = 4$ and $L = 100$, $C = 20$) and two empirical (*lac* promoter and GB1 protein). For each landscape, I fit $L^{-1} \log \rho = A + B \epsilon^a$ to the saddle-point density in the near-peak region, with A , B , and α as free parameters. The two simulated landscapes give $\alpha \approx 0.53$, close to the theoretical prediction of $\frac{1}{2}$ for i.i.d. fitness effects. The two empirical landscapes give $\alpha \approx 0.40$, slightly below $1/2$, consistent with a mild enrichment of small gaps. The same curves plotted against ϵ^a , using the fitted α value for each panel, are shown in SI Supplementary Fig. S2.

It is useful to ask over what fraction of the fitness range the near-peak scaling outperforms the Gaussian approximation. To quantify this, I define the *crossover fitness* F_{cross} as the fitness value at which ρ_{peak} becomes a closer approximation to ρ_{saddle} than ρ_{bulk} is. The crossover ratio $r = (F_{\text{max}} - F_{\text{cross}})/(F_{\text{max}} - F_{\text{mean}})$ measures the fraction of the fitness range (from F_{mean} to F_{max}) over which the near-peak scaling dominates. Because both $\log \rho_{\text{bulk}}$ and $\log \rho_{\text{peak}}$ scale linearly with L , the fitness at which they cross is determined by an L -independent equation, and so the crossover ratio r is independent of sequence length. Figure 4 marks F_{cross} for each of the four landscapes. Across these landscapes, r ranges from 0.26 to

0.34, indicating that the near-peak scaling governs roughly a quarter to a third of the fitness range.

Single-trait global epistasis models

The results above apply to purely additive landscapes. In many experimental settings, however, the observed fitness is better described not as a direct sum of per-position effects but rather as a nonlinear readout of an underlying additive trait. More precisely, the measured fitness takes the form $F = g(\phi)$, where $\phi(x) = \theta_0 + \sum_{l,c} \theta_{lc} x_{lc}$ is an additive trait and g is a nonlinear function. Such models are commonly referred to as global epistasis models (Starr and Thornton 2016; Sailer and Harms 2017; Otwinowski et al. 2018; Domingo et al. 2019). (Note that this definition of global epistasis is based on the structure of the underlying equations and is distinct from the phenomenological definition sometimes used in discussions of diminishing returns epistasis, e.g. Kryazhimskiy et al. 2014; Reddy and Desai 2021.) I now show that, when the fitness maximum occurs at an extremal value of the additive trait and the slope of g is nonzero there, the near-peak scaling exponent α is preserved, while only the coefficients A and B are modified.

The key observation is that, sufficiently close to the fitness peak, g is well-approximated by its linearization about ϕ_{max} . Under this linearization, the fitness deficit at each sequence is simply a




Fig. 5. Scaling under global epistasis. A simulated additive landscape ($L = 100$, $C = 4$, Gaussian θ_{lc}) for a trait ϕ is transformed by sigmoid nonlinearities $g(\phi)$ of varying steepness. (a) The log genotypic density $\log \rho_\phi$ in trait space versus ϕ , showing the saddle-point approximation (orange solid line) and the near-peak approximation (blue dashed line). Vertical dotted lines indicate $\phi_{\max}$ (red) and ϕ_{mean} (green). (b) The same quantities plotted against $\sqrt{\epsilon_F}$, showing the region of approximate linearity from which B_ϕ is extracted. (c) Three sigmoid nonlinearities with half-max location set to $\phi_{1/2} = \phi_{\max} - 4\sigma$ (gray dotted line) and varying span parameter s_ϕ , which controls steepness at half max. (d) The transformed density $\log \rho_F$ versus $\sqrt{\epsilon_F}$ for each sigmoid (solid lines) compared to the theoretical prediction $A_F + B_F \sqrt{\epsilon_F}$ (dashed lines in matching colors), where A_F and B_F are computed analytically from A_ϕ , B_ϕ , and $g'_{\max}$ with no additional fitting (see SI Sec. S5).

rescaled version of its trait deficit: $F_{\max} - F \approx (\phi_{\max} - \phi)g'_{\max}$, where $g'_{\max} = g'(\phi_{\max})$. Because this rescaling is linear, it preserves the power-law form of the gap distribution near zero and therefore preserves the scaling exponent α .

To state this more precisely, recall that the near-peak scaling hypothesis for the genotypic density of an additive trait (Equation (21)) is

$$\frac{1}{L} \log \rho_\phi \approx A_\phi + B_\phi \epsilon_\phi^\alpha, \quad (22)$$

where $\epsilon_\phi = (\phi_{\max} - \phi)/L$ is the per-site trait deficit. Applying the linearization to transform from trait space to fitness space, and defining the per-site fitness deficit $\epsilon_F = (F_{\max} - F)/L$, gives (see SI Sec. S5 for derivation):

$$\frac{1}{L} \log \rho_F \approx A_F + B_F \epsilon_F^\alpha, \quad (23)$$

where $A_F = A_\phi - \frac{1}{L} \log g'_{\max}$ and $B_F = (g'_{\max})^{-\alpha} B_\phi$. This approximation holds as long as the linearization of g remains accurate over the fitness range of interest. In practice, the flatter the nonlinearity is at $\phi_{\max}$, the more heavily it compresses the range over which the near-peak scaling hypothesis applies.

Figure 5 illustrates these predictions on a simulated landscape ($L = 100$, $C = 4$) passed through sigmoid nonlinearities of varying

steepness. The coefficients A_F and B_F are computed analytically from the scaling behavior of the additive trait ϕ and from the value of $g'_{\max}$, with no additional fitting. In all cases the predicted $\sqrt{\epsilon_F}$ scaling closely matches the exact transformed density, confirming that the near-peak exponent is preserved under global epistasis. As expected, sigmoids that are flatter at $\phi_{\max}$ (due to greater saturation) yield a larger B_F and a more compressed range of validity.

Multi-trait global epistasis models

The preceding section considered fitness landscapes defined by a nonlinear function of a single additive trait. More generally, however, fitness may depend on multiple additive traits simultaneously. I now show that the saddle-point machinery and the near-peak scaling results derived above extend naturally to this multi-trait setting. Consider K additive traits

$$\phi_k(x) = \theta_{0k} + \sum_{l,c} \theta_{lck} x_{lc}, \quad k = 1, \dots, K, \quad (24)$$

Taken together, these map each sequence x to a K -dimensional trait vector $\vec{\phi}(x) = (\phi_1, \dots, \phi_K)$. To define a tilted distribution, we introduce a complementary vector of K inverse temperatures $\vec{\beta} = (\beta_1, \dots, \beta_K)$. The resulting tilted distribution has the form

$p_{\vec{\beta}}(x) \propto e^{\vec{\beta} \cdot \vec{\phi}(x)}$. The key observation is that the exponent $\vec{\beta} \cdot \vec{\phi}(x)$ is itself an additive function of x , with effective per-position parameters $\tilde{\theta}_{ic} = \sum_k \beta_k \theta_{ick}$. The tilted distribution therefore factorizes across positions exactly as in the single-trait case. The corresponding cumulant-generating function is

$$\Phi(\vec{\beta}) = \sum_i \log \sum_c \exp \left(\sum_k \beta_k \theta_{ick} \right). \quad (20)$$

Its gradient $\nabla \Phi$ gives the mean of each trait under the tilted distribution, and the entries of the Hessian matrix Φ'' give the covariances.

The saddle-point approximation generalizes directly (Barndorff-Nielsen and Cox 1989, Sec. 6.5). As in the single-trait case, we choose $\vec{\beta}$ so that the target trait vector $\vec{\phi}$ coincides with the mean of the tilted distribution. Because Φ is strictly convex, this condition uniquely determines $\vec{\beta}$. The resulting genotypic density, which is joint across K traits, is

$$\rho_{\text{saddle}}(\vec{\phi}) = \frac{e^{\Phi(\vec{\beta}) - \vec{\beta} \cdot \vec{\phi}}}{(2\pi)^{K/2} |\Phi''(\vec{\beta})|^{1/2}}, \quad \text{where } \nabla \Phi(\vec{\beta}) = \vec{\phi}, \quad (26)$$

and where $|\Phi''|$ denotes the determinant of the Hessian. Note that this reduces to Equation (11) when $K = 1$.

The near-peak scaling behavior of single-trait global epistasis models also extends to the multi-trait setting. Suppose fitness is a nonlinear function of the trait vector, $F = g(\vec{\phi})$. The achievable trait vectors form a bounded region in K -dimensional trait space. If g attains its maximum F_{max} on the boundary of this region and has nonzero gradient there, then g can be linearized in the vicinity of the maximum. This linearization projects the K traits onto a single effective additive trait, $\tilde{\phi} = \sum_k c_k \phi_k$, where $\vec{c} = \nabla g(\vec{\phi}_{\text{max}})$. The near-peak scaling results from the preceding sections then apply directly, with the exponent α determined by the gap distribution of $\tilde{\phi}$. As in the single-trait case, the range of fitness values over which this scaling holds is curtailed by the nonlinearity of g .

Discussion

The core finding of this work is that the log genotypic density of additive fitness landscapes typically follows a power law near maximal fitness (Equation (21)). The exponent in the power law, α , determines the shape of this density near the peak: α is constrained to lie between 0 and 1, with larger α producing sharper peaks and smaller α producing stubbier ones. The value of α depends on how the fitness gaps between optimal and runner-up characters at each position are distributed, and thus varies from landscape to landscape. Notably, α is independent of both sequence length and alphabet size. In idealized landscapes, where the additive effects are drawn from a single distribution, one obtains $\alpha = \frac{1}{2}$ regardless of the specific distribution used. In more realistic situations, however, the scale of additive effects can vary substantially from position to position, causing α to deviate from $\frac{1}{2}$. In particular, one should often expect to find $\alpha < \frac{1}{2}$ (increased stubbiness) on empirical landscapes because positions with small-to-negligible effects are often much more frequent than positions with large effects.

These results, to my knowledge, provide the first quantitative description of near-peak abundance of genotypes for any class of fitness landscapes (apart from the highly idealized Hamming-distance landscape studied by Berg and von Hippel 1987). Indeed, the Gaussian approximation suggested by the CLT fundamentally breaks down near the fitness maximum of

additive landscapes. If taken literally, it predicts that many sequences lie far above F_{max} . If the Gaussian approximation is instead truncated at maximal fitness, it still predicts a vast overabundance of high-fitness sequences, often yielding densities that are off by many orders of magnitude. By contrast, the exponentiated power law that I describe is highly accurate in this regime. It reveals that, as fitness is decreased from its maximum, the number of available genotypes initially expands at a near-infinite rate, then grows increasingly slowly until it settles down to the growth rate predicted by the Gaussian approximation. The result is a broad and gently rounded (i.e. stubby) summit.

The exponentiated power law given in Equation (21) comes with several caveats. First, it applies only to additive and globally epistatic landscapes. In particular, the consequences of specific epistatic interactions on near-peak density have yet to be investigated. Second, the derivation assumes that sequence length L is large. In practice, however, this scaling behavior is often observed even for landscapes describing sequences of length $L \sim 10$, such as TF binding motifs; see Supplementary Fig. S3. Third, the derivation assumes that the distribution of position-specific fitness gaps behaves as a power law near zero. That said, Equation (21) appears to provide a good empirical fit to many additive landscapes, even when α , A , and B must be fit empirically rather than derived from the gap distribution.

The scaling behavior also extends to globally epistatic landscapes defined on one or more additive traits. This only happens, however, under certain conditions: the nonlinearity must attain its maximum value on the boundary of the accessible trait space, and it must do so with nonzero gradient. If these conditions are satisfied, the nonlinearity can be approximated by a linear function near the peak, and the problem of computing genotypic densities reduces to that of a one-dimensional additive fitness function. If there is just one additive trait, the near-peak density will exhibit the same exponent α (and thus the same “stubbiness”) as the additive trait itself. When there are multiple additive traits, the linearized fitness landscape will be a linear combination of the traits, and is therefore expected to exhibit an empirical α intermediate to the α values for the individual traits (the math here is not so clean, though). An important caveat with all such global epistasis models is that the scaling behavior will only be valid over the range of F values for which linearization provides a good approximation. This range will be particularly small for global epistasis nonlinearities that plateau near their maxima.

Looking forward, the mathematical methods used in this work might prove useful for addressing other questions about fitness landscapes. The distribution of fitness effects (DFE) is of central importance in evolutionary genetics (Gillespie 1983; Eyre-Walker and Keightley 2007), and future work using related mathematical approaches could potentially be used to study the DFE near fitness peaks. Related questions, such as the distribution of Hamming distances among sequences at a given fitness level, might also be addressed. Finally, the multi-trait saddle-point approximation in Equation (26) naturally suggests applications to Fisher's geometric model (Fisher 1930; Orr 1998). It would be particularly interesting, in the spirit of Hwang et al. (2017), to see if similar techniques could shed light on how the dramatic variation in genotypic density that occurs across the fitness range affects the predictions of Fisher's model.

Data availability

All the code used to carry out this analysis and generate the figures is publicly available at https://github.com/jbkinney/26_stubby.

A snapshot of this repository is available on Zenodo at <https://doi.org/10.5281/zenodo.20402692>. Supplemental Information is available online at [GENETICS](https://www.genetics.org/genetics).

Acknowledgments

This project grew out of conversations with Ville Mustonen and Michael Lässig back in 2006. I also thank my colleagues David McCandlish and Peter Koo for helpful discussions.

Funding

Financial support was provided by the US National Institutes of Health, grants R01HG011787 and S10OD028632.

Conflicts of interest

None declared.

Literature cited

- Adams RM, Mora T, Walczak AM, Kinney JB. 2016. Measuring the sequence-affinity landscape of antibodies with massively parallel titration curves. *Elife*. 5:e23156. <https://doi.org/10.7554/eLife.23156>.
- Aita T, Husimi Y. 1996. Fitness spectrum among random mutants on Mt. Fuji-type fitness landscape. *J Theor Biol*. 182:469–485. <https://doi.org/10.1006/jtbi.1996.0189>.
- Aita T, Husimi Y. 2000. Adaptive walks by the fittest among finite random mutants on a Mt. Fuji-type fitness landscape II. effect of small non-additivity. *J Math Biol*. 41:207–231. <https://doi.org/10.1007/s002850000046>.
- Aita T, Morinaga S, Husimi Y. 2004. Thermodynamical interpretation of evolutionary dynamics on a fitness landscape in a evolution reactor, I. *Bull Math Biol*. 66:1371–1403. <https://doi.org/10.1016/j.bulm.2004.01.004>.
- Baeza-Centurion P, Miñana B, Schmiedel JM, Valcárcel J, Lehner B. 2019. Combinatorial genetics reveals a scaling law for the effects of mutations on splicing. *Cell*. 176:549–563.e23. <https://doi.org/10.1016/j.cell.2018.12.010>.
- Barndorff-Nielsen O, Cox D. 1989. Asymptotic techniques for use in statistics. Monographs on statistics and applied probability. Springer US.
- Barton NH, Coe JB. 2009. On the application of statistical physics to evolutionary biology. *J Theor Biol*. 259:317–324. <https://doi.org/10.1016/j.jtbi.2009.03.019>.
- Barton NH, de Vladar HP. 2009. Statistical mechanics and the evolution of polygenic quantitative traits. *Genetics*. 181:997–1011. <https://doi.org/10.1534/genetics.108.099309>.
- Barton NH, Etheridge AM, Véber A. 2017. The infinitesimal model: definition, derivation, and implications. *Theor Popul Biol*. 118: 50–73. <https://doi.org/10.1016/j.tpb.2017.06.001>.
- Barton NH, Keightley PD. 2002. Understanding quantitative genetic variation. *Nat Rev Genet*. 3:11–21. <https://doi.org/10.1038/nrg700>.
- Baydar DO et al. 2025. JASPAR 2026: expansion of transcription factor binding profiles and integration of deep learning models. *Nucleic Acids Res*. 54:D184–D193. <https://doi.org/10.1093/nar/gkaf1209>.
- Berg J, Willmann S, Lässig M. 2004. Adaptive evolution of transcription factor binding sites. *BMC Evol Biol*. 4:42. <https://doi.org/10.1186/1471-2148-4-42>.
- Berg OG, von Hippel PH. 1987. Selection of DNA binding sites by regulatory proteins. Statistical-mechanical theory and application to operators and promoters. *J Mol Biol*. 193:723–743. [https://doi.org/10.1016/0022-2836\(87\)90354-8](https://doi.org/10.1016/0022-2836(87)90354-8).
- Berg OG, von Hippel PH. 1988. Selection of DNA binding sites by regulatory proteins. II. The binding specificity of cyclic AMP receptor protein to recognition sites. *J Mol Biol*. 200:709–723. [https://doi.org/10.1016/0022-2836\(88\)90482-2](https://doi.org/10.1016/0022-2836(88)90482-2).
- Carneiro M, Hartl DL. 2010. Adaptive landscapes and protein evolution. *Proc Natl Acad Sci U S A*. 107:1747–1751. <https://doi.org/10.1073/pnas.0906192106>.
- Daniels HE. 1954. Saddlepoint approximations in statistics. *Ann Math Stat*. 25:631–650. <https://doi.org/10.1214/aoms/1177728652>.
- Darwin CG, Fowler RH. 1922. On the partition of energy. *Philos Mag Ser 6*. 44:450–479. <https://doi.org/10.1080/14786440908565189>.
- Djordjevic M, Sengupta AM, Shraiman BI. 2003. A biophysical approach to transcription factor binding site discovery. *Genome Res*. 13:2381–2390. <https://doi.org/10.1101/gr.1271603>.
- Domingo J, Baeza-Centurion P, Lehner B. 2019. The causes and consequences of genetic interactions (epistasis). *Annu Rev Genomics Hum Genet*. 20:433–460. <https://doi.org/10.1146/annurev-genom-083118-014857>.
- Elemento O, Slonim N, Tavazoie S. 2007. A universal framework for regulatory element discovery across all genomes and data types. *Mol Cell*. 28:337–350. <https://doi.org/10.1016/j.molcel.2007.09.027>.
- Eyre-Walker A, Keightley PD. 2007. The distribution of fitness effects of new mutations. *Nat Rev Genet*. 8:610–618. <https://doi.org/10.1038/nrg2146>.
- Faure AJ et al. 2024. The genetic architecture of protein stability. *Nature*. 634:995–1003. <https://doi.org/10.1038/s41586-024-07966-0>.
- Fisher RA. 1930. The genetical theory of natural selection. Clarendon Press.
- Foat B, Morozov A, Bussemaker H. 2006. Statistical mechanical modeling of genome-wide transcription factor occupancy data by MatrixREDUCE. *Bioinformatics*. 22:e141–e149. <https://doi.org/10.1093/bioinformatics/btl223>.
- Fragata I, Blanckaert A, Dias Louro MA, Liberles DA, Bank C. 2019. Evolution in the light of fitness landscape theory. *Trends Ecol Evol*. 34:69–82. <https://doi.org/10.1016/j.tree.2018.10.009>.
- Gerland U, Moroz JD, Hwa T. 2002. Physical constraints and functional characteristics of transcription factor–DNA interaction. *Proc Natl Acad Sci U S A*. 99:12015–12020. <https://doi.org/10.1073/pnas.192693599>.
- Gillespie JH. 1983. A simple stochastic gene substitution model. *Theor Popul Biol*. 23:202–215. [https://doi.org/10.1016/0040-5809\(83\)90014-x](https://doi.org/10.1016/0040-5809(83)90014-x).
- Grant CE, Bailey TL, Noble WS. 2011. FIMO: scanning for occurrences of a given motif. *Bioinformatics*. 27:1017–1018. <https://doi.org/10.1093/bioinformatics/btr064>.
- Hill WG, Goddard ME, Visscher PM. 2008. Data and theory point to mainly additive genetic variance for complex traits. *PLoS Genet*. 4:e1000008. <https://doi.org/10.1371/journal.pgen.1000008>.
- Hwang S, Park SC, Krug J. 2017. Genotypic complexity of Fisher's geometric model. *Genetics*. 206:1049–1079. <https://doi.org/10.1534/genetics.116.199497>.
- Iwasa Y. 1988. Free fitness that always increases in evolution. *J Theor Biol*. 135:265–281. [https://doi.org/10.1016/s0022-5193\(88\)80243-1](https://doi.org/10.1016/s0022-5193(88)80243-1).
- Kinney JB, Murugan A, Callan CG, Cox EC. 2010. Using deep sequencing to characterize the biophysical mechanism of a transcriptional regulatory sequence. *Proc Natl Acad Sci U S A*. 107: 9158–9163. <https://doi.org/10.1073/pnas.1004290107>.
- Kinney JB, Tkacik G, Callan CG. 2007. Precise physical models of protein–DNA interaction from high-throughput data. *Proc Natl Acad Sci U S A*. 104:501–506. <https://doi.org/10.1073/pnas.0609908104>.

- Kryazhimskiy S, Rice DP, Jerison ER, Desai MM. 2014. Global epistasis makes adaptation predictable despite sequence-level stochasticity. *Science*. 344:1519–1522. <https://doi.org/10.1126/science.1250939>.
- Lässig M. 2007. From biophysics to evolutionary genetics: statistical aspects of gene regulation. *BMC Bioinformatics*. 8:S7. <https://doi.org/10.1186/1471-2105-8-S6-S7>.
- Madsen T, Hobolth A, Jensen JL, Pedersen JS. 2017. Significance evaluation in factor graphs. *BMC Bioinformatics*. 18:199. <https://doi.org/10.1186/s12859-017-1614-z>.
- Maynard Smith J. 1970. Natural selection and the concept of a protein space. *Nature*. 225:563–564. <https://doi.org/10.1038/225563a0>.
- Mustonen V, Kinney J, Callan CG, Lässig M. 2008. Energy-dependent fitness: a quantitative model for the evolution of yeast transcription factor binding sites. *Proc Natl Acad Sci U S A*. 105:12376–12381. <https://doi.org/10.1073/pnas.0805909105>.
- Mustonen V, Lässig M. 2005. Evolutionary population genetics of promoters: predicting binding sites and functional phylogenies. *Proc Natl Acad Sci U S A*. 102:15936–15941. <https://doi.org/10.1073/pnas.0505537102>.
- Neher RA, Shraiman BI. 2011. Statistical genetics and evolution of quantitative traits. *Rev Mod Phys*. 83:1283–1300. <https://doi.org/10.1103/RevModPhys.83.1283>.
- Neidhart J, Szendro IG, Krug J. 2014. Adaptation in tunably rugged fitness landscapes: the rough Mount Fuji model. *Genetics*. 198: 699–721. <https://doi.org/10.1534/genetics.114.167668>.
- Olson CA, Wu NC, Sun R. 2014. A comprehensive biophysical description of pairwise epistasis throughout an entire protein domain. *Curr Biol*. 24:2643–2651. <https://doi.org/10.1016/j.cub.2014.09.072>.
- Orr HA. 1998. The population genetics of adaptation: the distribution of factors fixed during adaptive evolution. *Evolution*. 52:935–949. <https://doi.org/10.1111/j.1558-5646.1998.tb01823.x>.
- Otwinowski J. 2018. Biophysical inference of epistasis and the effects of mutations on protein stability and function. *Mol Biol Evol*. 35: 2345–2354. <https://doi.org/10.1093/molbev/msy141>.
- Otwinowski J, McCandlish DM, Plotkin JB. 2018. Inferring the shape of global epistasis. *Proc Natl Acad Sci U S A*. 115:E7550–E7558. <https://doi.org/10.1073/pnas.1804015115>.
- Pachkov M, Erb I, Molina N, Nimwegen E. 2007. SwissRegulon: a database of genome-wide annotations of regulatory sites. *Nucleic Acids Res*. 35:D127–D131. <https://doi.org/10.1093/nar/gkl857>.
- Pahujani S, Krug J. 2025. Complexity and accessibility of random landscapes. *SciPost Phys Lect Notes*. 108:1–23. <https://doi.org/10.21468/SciPostPhysLectNotes.108>.
- Reddy G, Desai MM. 2021. Global epistasis emerges from a generic model of a complex trait. *Elife*. 10:e64740. <https://doi.org/10.7554/eLife.64740>.
- Reid N. 1988. Saddlepoint methods and statistical inference. *Stat Sci*. 3:213–238. <https://doi.org/10.1214/ss/1177012906>.
- Sailer ZR, Harms MJ. 2017. Detecting high-order epistasis in non-linear genotype-phenotype maps. *Genetics*. 205:1079–1088. <https://doi.org/10.1534/genetics.116.195214>.
- Sarkisyan KS et al. 2016. Local fitness landscape of the green fluorescent protein. *Nature*. 533:397–401. <https://doi.org/10.1038/nature17995>.
- Sella G, Hirsh AE. 2005. The application of statistical physics to evolutionary biology. *Proc Natl Acad Sci U S A*. 102:9541–9546. <https://doi.org/10.1073/pnas.0501865102>.
- Sengupta AM, Djordjevic M, Shraiman BI. 2002. Specificity and robustness in transcription control networks. *Proc Natl Acad Sci U S A*. 99:2072–2077. <https://doi.org/10.1073/pnas.022388499>.
- Staden R. 1989. Methods for calculating the probabilities of finding patterns in sequences. *Comput Appl Biosci*. 5:89–96. <https://doi.org/10.1093/bioinformatics/5.2.89>.
- Starr TN et al., 2020. Deep mutational scanning of SARS-CoV-2 receptor binding domain reveals constraints on folding and ACE2 binding. *Cell*. 182:1295–1310.e20. <https://doi.org/10.1016/j.cell.2020.08.012>.
- Starr TN, Thornton JW. 2016. Epistasis in protein evolution. *Protein Sci*. 25:1204–1218. <https://doi.org/10.1002/pro.2897>.
- Stiffler MA, Hekstra DR, Ranganathan R. 2015. Evolvability as a function of purifying selection in TEM-1 β -lactamase. *Cell*. 160: 882–892. <https://doi.org/10.1016/j.cell.2015.01.035>.
- Stormo G, Fields D. 1998. Specificity, free energy and information content in protein-DNA interactions. *Trends Biochem Sci*. 23: 109–113. [https://doi.org/10.1016/s0968-0004\(98\)01187-6](https://doi.org/10.1016/s0968-0004(98)01187-6).
- Stormo GD. 2000. DNA binding sites: representation and discovery. *Bioinformatics*. 16:16–23. <https://doi.org/10.1093/bioinformatics/16.1.16>.
- Szendro IG, Schenk MF, Franke J, Krug J, de Visser JAGM. 2013. Quantitative analyses of empirical fitness landscapes. *J Stat Mech*. 2013:P01005. <https://doi.org/10.1088/1742-5468/2013/01/P01005>.
- Tareen A et al. 2022. MAVE-NN: learning genotype-phenotype maps from multiplex assays of variant effect. *Genome Biol*. 23:98. <https://doi.org/10.1186/s13059-022-02661-7>.
- Tareen A, Kinney JB. 2019. Logomaker: beautiful sequence logos in Python. *Bioinformatics*. 36:2272–2274. <https://doi.org/10.1093/bioinformatics/btz921>.
- Touchette H. 2009. The large deviation approach to statistical mechanics. *Phys Rep*. 478:1–69. <https://doi.org/10.1016/j.physrep.2009.05.002>.
- Touzet H, Varré JS. 2007. Efficient and accurate P-value computation for Position Weight Matrices. *Algorithms Mol Biol*. 2:15. <https://doi.org/10.1186/1748-7188-2-15>.
- Turelli M, Barton NH. 1994. Genetic and statistical analyses of strong selection on polygenic traits: what, me normal? *Genetics*. 138: 913–941. <https://doi.org/10.1093/genetics/138.3.913>.
- Vorontsov IE et al. 2024. HOCOMOCO in 2024: a rebuild of the curated collection of binding models for human and mouse transcription factors. *Nucleic Acids Res*. 52:D154–D163. <https://doi.org/10.1093/nar/gkad1077>.
- Walsh B, Lynch M. 2018. Evolution and selection of quantitative traits. Oxford University Press.
- Wingender E et al. 2000. TRANSFAC: an integrated system for gene expression regulation. *Nucleic Acids Res*. 28:316–319. <https://doi.org/10.1093/nar/28.1.316>.
- Wright S. 1932. The roles of mutation, inbreeding, crossbreeding, and selection in evolution. In: *Proceedings of the Sixth International Congress of Genetics*. 1. p. 356–366.

Editor: J. Krug